

Leah Cunningham, Karl Eichwalder, Thomas Fehr, Stefan Fent, Werner Fink, Peter Findeisen, Dennis Geider, Viviane Glanz, Carsten Groß, Roland Haidl, Björn Jacke, Richard Jelinek, Hubert Mantel, Johannes Meixner, Edith Parzefall, Peter Pöml, Peter Reinhart, Marc Rührschneck, Thomas Schraitle, Klaus G. Wagner, Christian Zoz

SuSE Linux


System and Reference Manual

SuSE Inc.
580 2nd Street, #210
Oakland, CA 94607
USA
Toll free phone number
within the US and Canada: 1-888-UR-LINUX (1-888-875-4689)
Phone.: +1-510-628-3380
Fax.: +1-510-628-3381
E-mail: info@suse.com
orders@suse.com
webmaster@suse.com
WWW: <http://www.suse.com>

Europe:

SuSE Linux Ltd.
The Kinetic Centre
Theobald Street
Borehamwood, WD6 4PJ UK
Phone.: +44-20-8387-4088
Fax: +44-20-8387-4010
<http://www.suse.co.uk>

SuSE GmbH
Schanzäckerstr. 10
D-90443 Nürnberg
Germany
Tel.: +49-911-740-5331
Fax.: +49-911-741 77 55
E-mail: suse@suse.de
WWW: <http://www.suse.de>



Leah Cunningham, Karl Eichwalder, Thomas Fehr, Stefan Fent, Werner Fink, Peter Findeisen, Dennis Geider, Viviane Glanz, Carsten Groß, Roland Haidl, Björn Jacke, Richard Jelinek, Hubert Mantel, Johannes Meixner, Edith Parzefall, Peter Pöml, Peter Reinhart, Marc Rührschneck, Thomas Schraitle, Klaus G. Wagner, Christian Zoz

SuSE Linux
System and Reference Manual
2nd revised edition 2001
SuSE GmbH

Copyright

This work is copyrighted by SuSE GmbH.

You may copy it in whole or in part as long as the copies retain this copyright statement.

Layout: \LaTeX

English Translation and Editing: Rebecca Ellis, Marc Rührschneck

Design: Manuela Piotrowski

Graphics: Rolf Vogt

Linux is a registered trademark of *Linus Torvalds*. *XFree86*™ is a registered trademark of *The XFree86 Project, Inc.* *MS-DOS*, *Windows*, *Windows 95*, *Windows 98*, and *Windows NT* are registered trademarks of *Microsoft Corporation*. *UNIX* is a registered trademark of *X/Open Company Limited*. Other trademarks and registered trademarks are: *Clipper* of *Computer Associates*, *dBASE* of *Borland*, *Foxbase* of *Microsoft*, *Compuserve*, *T-Online* of *Deutsche Telekom*, *SuSE* and *YaST* of *SuSE GmbH*. All trade names are used without the guarantee for their free use and are possibly registered trade marks. SuSE GmbH essentially follows the notations of the manufacturers. Other products mentioned in this manual may be trademarks of the respective manufacturer.

Contents

1	Introduction	1
1.1	What is “Linux”?	1
1.2	Purpose of this Book — Notes on Usage	1
1.3	Typographical Conventions — What Does “earth:~ # ls ” Mean	3
1.4	Acknowledgments	3
2	Your First SuSE Linux Installation	7
2.1	Installation with the Text-Based YaST (YaST 1)	7
2.1.1	Starting: The Welcome Screen	7
2.1.2	The Basics: linuxrc	8
2.1.3	Installation with YaST	12
2.1.4	Partitioning and Formatting the Hard Disk	13
2.1.5	Installation of Software Packages	15
2.1.6	Selection of the Base Software	16
2.1.7	Installing System Software and Programs	17
2.1.8	Selecting a Kernel for the System	18
2.1.9	Base System Configuration with YaST	19
2.1.10	Logging in After Your First Installation	22
2.1.11	Shutting Down and Booting the System	23
2.2	Starting SuSE Linux	24
2.3	Special Installations	26
2.3.1	Installation Without a Supported CD-ROM Drive	26
2.3.2	Automatic Installation and Configuration of SuSE Linux	27
2.4	Tips and Tricks	28
2.4.1	Creating Space for Linux (Partitioning)	28
2.4.2	Booting from Disk (SYSLINUX)	31
2.4.3	Booting with CD 2	31
2.4.4	Creating a Boot Disk In DOS	32
2.4.5	Creating a Boot Disk with UNIX	33
2.4.6	Does Linux Support my CD-ROM?	34
2.4.7	ATAPI CD-ROM Hangs While Reading	34
2.4.8	Problems with CD-ROM Drives on Parallel Port	35

2.4.9	Loadlin Does Not Have Enough Memory to Load the Kernel	36
2.4.10	Loadlin Does Not Start	36
2.5	Partitioning Tips	36
2.6	Partitioning for Experts	39
2.6.1	Size of Swap Partition	39
2.6.2	Computer Used as Stand-Alone Machine	39
2.6.3	Optimizations	41
3	YaST — Yet another Setup Tool	43
3.1	Basic Usage and Keyboard Layout	43
3.2	The YaST Main Menu	43
3.3	General Help	44
3.4	Adjusting Installation	45
3.4.1	Selecting the Language	45
3.4.2	Selecting Keyboard Layout	45
3.4.3	Installation Medium	45
3.4.4	Installation from CD-ROM or DVD	46
3.4.5	Installation from a Hard Disk Partition	46
3.4.6	Installation via NFS	47
3.4.7	Installation from a Reachable Directory	48
3.4.8	Installation via FTP	49
3.4.9	Partitioning Your Hard Disk	50
3.4.10	Setting Target Partitions and File Systems	52
3.4.11	Configuring Your Hard Disk Manually	56
3.4.12	File systems and Mount Points	58
3.4.13	Configuring the Logical Volume Manager	60
3.4.14	Installation to a Directory	61
3.5	Choosing and Installing Packages	61
3.5.1	Load Configuration	62
3.5.2	Save Configuration	62
3.5.3	Changing Your Configuration	63
3.5.4	Checking Package Dependencies	65
3.5.5	What If...	67
3.5.6	Start Installation	67
3.5.7	Index of All Series and Packages	67
3.5.8	Searching for Packages	67
3.5.9	Installing Packages	67
3.5.10	Deleting Packages	69

3.6	Updating the System	69
3.7	System Administration	69
3.7.1	Integrating Hardware into the System	70
3.7.2	Kernel and Boot Configuration	71
3.7.3	Network Configuration	74
3.7.4	Login Configuration	76
3.7.5	Settings for susewm (the Window Manager)	77
3.7.6	User Administration	78
3.7.7	Group Administration	79
3.7.8	Creating Backups	80
3.7.9	System Security Settings	81
3.7.10	Setting the Console Font	81
3.7.11	Setting the time zone	81
3.7.12	Configuring XFree86(TM)	81
3.7.13	Configuring gpm	82
3.7.14	Changing the Configuration File	82
4	YaST2 — Configuration Tools	85
4.1	Hardware	86
4.2	Network/Basic	86
4.3	Network/Advanced	86
4.4	Security and Users	86
4.5	Software	86
4.5.1	Patch CD Update	87
4.6	System	87
4.6.1	RC.Config Editor	87
4.6.2	Runlevel Editor	87
4.6.3	Expert Partitioner	88
4.6.4	Logical Volume Manager (LVM)	89
4.6.5	Soft RAID	94
4.6.6	Kernel Selection	96
4.7	Miscellaneous	96
4.7.1	Printers for CUPS and LPD	97
5	Booting and Boot Managers	99
5.1	Booting a PC	99
5.2	Boot Concepts	100
5.3	An Overview of LILO	101
5.4	Configuring LILO	104

5.4.1	Structure of <code>lilo.conf</code>	105
5.4.2	Other LILO Configuration Options	107
5.5	Installing and Uninstalling LILO	110
5.6	Creating a Linux Boot Disk	112
5.7	Sample Configurations	114
5.7.1	DOS/Windows 95/98 and Linux	114
5.7.2	Windows NT and Linux on One Hard Disk	115
5.7.3	OS/2 and Linux	116
5.7.4	DOS, OS/2, and Linux	116
5.8	LILO Problems	117
5.8.1	Diagnosis of Errors: LILO Start Messages	118
5.8.2	The 1024-Cylinder Limit	119
5.8.3	Special Boot Problems with Kernels from 2.0 Onwards	121
5.9	Starting via <code>loadlin</code>	122
5.9.1	Necessary Steps for All <code>loadlin</code> Users	123
5.9.2	Setting up Boot Menus	123
5.9.3	Starting Linux from Within Windows	125
5.9.4	The Windows Boot Menu	126
6	The X Window System	129
6.1	Historical Background	129
6.2	Version 4.x of XFree86	130
6.3	Configuration Using <code>SaX2</code>	132
6.3.1	Troubleshooting	132
6.3.2	The X Server Log File	136
6.3.3	Starting the X Window System	138
6.4	Configuration Using <code>SaX</code>	138
6.4.1	Reconfiguring	146
6.4.2	Troubleshooting	146
6.5	Configuration Using <code>xf86config</code>	147
6.6	Optimizing the Installation of the X Window System	154
6.6.1	Inserting Additional (True Type) Fonts	159
6.6.2	Setting Up Input Devices	161
6.6.3	3D Acceleration	163

7	Linux and Hardware	165
7.1	Preliminary Notes	165
7.2	Cards	165
7.2.1	PCI and ISA Cards	165
7.3	Sound Cards	169
7.3.1	Configuring Sound Cards with YaST2	169
7.3.2	Configuring Your Sound Card With alsacofn	169
7.4	Ports on a Computer	171
7.4.1	PS/2 Ports	171
7.4.2	Serial Ports	171
7.4.3	Parallel Ports	171
7.4.4	USB – Universal Serial Bus	173
7.5	Removable Drives	176
7.5.1	Floppy Disk Drives	176
7.5.2	LS-120 Drives	177
7.5.3	ZIP Drives	177
7.6	Modems	178
7.6.1	External Modems	178
7.6.2	Internal Modems	178
7.7	Tape Drives	180
8	Notebooks — PCMCIA, APM, IrDA	183
8.1	PCMCIA	183
8.1.1	Hardware	183
8.1.2	Software	184
8.1.3	Configuration	184
8.1.4	Configurations for Changing — “Schemes”	186
8.1.5	If Things Still Don’t Work	188
8.1.6	Installation via PCMCIA	190
8.1.7	Other Help Programs	191
8.1.8	Compiling PCMCIA	191
8.2	APM — Power Management	192
8.2.1	Fundamentals	192
8.2.2	The Correct Kernel	193
8.2.3	The APM Daemon	193
8.2.4	More Commands	194
8.2.5	Pause for the Hard Disk	194
8.3	IrDA – Infrared Data Association	194

9	Printing	197
9.1	Basics and Requirements for Printing	197
9.1.1	Different Standard Printer Languages	197
9.1.2	The Problem with GDI Printers	197
9.1.3	Hardware and Software Requirements	198
9.1.4	Defining a Suitable Printer Driver	199
9.2	Connecting the Printer Locally	199
9.2.1	Parallel Ports	199
9.2.2	Special Slot Cards: ISA-PnP and PCI	201
9.2.3	USB Connections	202
9.2.4	Serial Ports	203
9.3	lpd: the Print Manager and Tools	204
9.3.1	The Print Manager in the Background	204
9.3.2	Additional Tools and Control Tools for the User	205
9.4	lpr: Sending Print Jobs	207
9.5	SETUP (lprsetup): Configuring the Print System	208
9.5.1	Hints on Configuring with lprsetup	208
9.5.2	Testing the Printer	210
9.5.3	Special Cases	210
9.6	The apsfiler	211
9.6.1	How the apsfiler Works	211
9.6.2	Configuring the apsfiler	212
9.6.3	Finding Errors with apsfiler	214
9.7	Printing in the TCP/IP Network	215
9.7.1	Overview: Specific Cases	215
9.7.2	Configuring a Printer Forwarding Queue on the Client	216
9.7.3	Configuring a Prefilter on a Client	217
9.7.4	Problems with Configuration	218
9.8	Some Words on Ghostscript	219
9.8.1	Examples of Working with Ghostscript	219
9.9	Examples of Your Own Print Filters	220
9.9.1	Preliminary Notes	220
9.9.2	A Simple Example of the Basic Method of Working	221
9.9.3	A More Complex Example	224

10 The Kernel	229
10.1 Kernel Sources	229
10.2 Kernel Modules	230
10.3 Kernel Configuration	232
10.4 Settings in the Kernel Configuration	233
10.5 Compiling the Kernel	233
10.6 Installing the Kernel	234
10.7 Creating a Boot Disk	235
10.8 Cleaning Your Hard Disk After Compilation	235
11 Kernel Parameters	237
11.1 Drivers in the Kernel	237
11.2 Some Tips	238
11.3 The Parameters	238
11.3.1 Notations and Meanings	238
11.3.2 Kernel Parameters at the Boot Prompt	239
11.3.3 CD-ROM Drives on Proprietary Controllers	249
11.3.4 <code>modprobe</code> Parameters	251
12 Updating the System and Package Management	271
12.1 Updating SuSE Linux	271
12.1.1 Preparations	271
12.1.2 Updating with YaST2	272
12.1.3 Updating with YaST1	272
12.1.4 Updating Single Packages	274
12.2 From Version to Version	274
12.2.1 From 6.4 to 7.0	275
12.2.2 From 7.0 to 7.1	276
12.2.3 From 7.1 to 7.2	276
12.2.4 From 7.2 to 7.3	277
12.3 RPM — the Distribution Package Manager	278
12.3.1 Verifying Package Authenticity	279
12.3.2 Managing Packages: Install, Update, and Uninstall	279
12.3.3 RPM Queries	280
12.3.4 Installing and Compiling Source Packages	282
12.3.5 Other Tools for Working with RPM Archives	283

13	Special Features of SuSE Linux	285
13.1	File System Hierarchy Standard (FHS) and the Linux Standard Base (LSB)	285
13.1.1	File System Hierarchy Standard (FHS)	285
13.1.2	Linux Standard Base (LSB)	285
13.1.3	teTeX — T _E X in SuSE Linux	285
13.2	Example Environments for FTP and HTTP	285
13.3	Hints on Special Software Packages	286
13.3.1	The packages <code>xdevel</code> and <code>xdevel33</code>	286
13.3.2	Package <code>cron</code>	286
13.3.3	Package <code>curses</code>	286
13.3.4	Sources for <code>uucp</code>	287
13.3.5	Man Pages	287
13.3.6	The Command <code>ulimit</code>	287
13.3.7	The File <code>/etc/resolv.conf</code>	288
13.4	Booting with the “initial ramdisk”	288
13.4.1	Concept of the Initial Ramdisk	288
13.4.2	The Order of the Booting Process with <code>initrd</code>	289
13.4.3	Boot Loaders	289
13.4.4	Using <code>initrd</code> in SuSE	290
13.4.5	Possible Difficulties — Self-Compiled Kernels	291
13.4.6	Prospects	292
13.5	<code>linuxrc</code>	292
13.6	The SuSE Rescue System	297
13.6.1	Preparations	297
13.6.2	Starting the Rescue System	298
13.6.3	Working with the Rescue System	299
13.7	Virtual Consoles	301
13.8	Keyboard Layout	301
13.9	Local Adjustments — I18N/L10N	303
13.9.1	Settings for Language Support	305
14	The SuSE Linux Boot Concept	307
14.1	The <code>init</code> Program	307
14.2	Runlevels	307
14.3	Changing Runlevels	309
14.4	Init Scripts	309
14.5	<code>/etc/rc.config</code> and <code>SuSEconfig</code>	312
14.6	The Variables in <code>/etc/rc.config</code> — System Configuration	313

A	Manual Page of e2fsck	325
B	The GNU General Public License	329
C	Important Keys	337
	Bibliography	339

1 Introduction

1.1 What is “Linux”?

The actual “Linux” is the *kernel*, the “heart” of every UNIX operating system. The kernel alone does not constitute an entire operating system, however. Many of the tools in UNIX are free software. Because of this, almost all standard tools in Unix are also available in Linux. It is these programs which make the operating system what it is.

Many of these tools are GNU versions of these Unix programs that offer enhanced features over the originals. Of these, the GNU C/C++ compiler, one of the best compilers available, is probably the most famous program. We should not forget the many tools, small and large, which can be used from the command line, in shell scripts, or as full fledged programs — from the shell, file and text utilities such as *sed*, *awk*, and *perl* to editors like *vi* or complete environments like Emacs.

What makes Linux complete is XFree86™ (current version, 4.1.0), the X Window System (currently X11 Release 6.4) for PC-based UNIX systems. This port is part of the official X11R6.4 distribution of the X Consortium, Inc., and is fully compatible with that standard.

All these components, together with many other tools, applications, and amusements, make up the system known as *Linux*. A large amount of free software is available and there are many ways to put a Linux system together. At this point, the Linux distributions come into play, in this case, SuSE Linux. A distribution is a selection from a wide palette of free software. The end product is, in the case of the SuSE Linux CDs, over 1000 software packages.

Note

SuSE Linux includes, in *addition* to free software — software in which the source code is included on the CD-ROMs, software packages that, for different reasons, SuSE can only offer in compiled form.



Because distributions are offered free on the Internet, updating your system does not necessarily mean buying a new distribution.

1.2 Purpose of this Book — Notes on Usage

This book should not and cannot be a replacement for more detailed literature. To dig deeper into the fascinating world of *high-performance computing* refer to

the bibliography at the end of the book (Chapter 14.6 page 339).

The dynamics of Linux and of free software is a big challenge as far as writing a handbook is concerned. It must keep up with the most recent developments and, at the same time, remain suited for those new to the Unix world without boring experienced Linux users. The book in front of you attempts to do justice to all these requirements. The aim of this book is to allow Linux beginners an encouraging start on a voyage of discovery through the new system.

The book is divided into the following sections:

Installation: This section guides the Linux novice with a documented example installation (Chapter 2 page 7). It points out possible sources of error and provides concrete help with problems.

YaST: The use of YaST, the SuSE installation and administration tool, is described in Chapter 3 page 43. Chapter 5 page 99 explains how to boot the kernel.

Graphical User Interface: Activating and setting up the X user interface is the subject of Chapter 6 page 129.

Sound, Scanner, Notebooks, and Printing: In Chapters 7 page 165 through 9 page 197, the capabilities of your hardware are explained.

The Linux Kernel: Here we get down to the heart of the matter. The Linux kernel is introduced in Chapter 10 page 229 and Chapter 11 page 237, providing a guide on compiling and installing your own kernel.

Update, Software Packages, Booting: Several update strategies as well as software management are introduced (Chapter 12 page 271), special features of SuSE Linux are covered, and the boot concept is explained (Chapter 14 page 307).

Special Features of SuSE Linux: Here you will learn more about the file system and the \TeX environment of SuSE Linux. Special software packages and the SuSE rescue system will also be discussed, among other things.

The SuSE Linux Boot Concept: In Chapter 14 page 307, read details on the init program, the various runlevels, and the system configuration files.

Technical Appendix In the appendix, find a list of the most important configuration files, frequently used key combinations, example files, a licensing agreement, and much more.

Bibliography and Index If there is something you still have not found, consult the extensive bibliography. Either find the place which explains important concepts or simply increase your general knowledge about Linux, Unix, or computers.

1.3 Typographical Conventions — What Does “earth:~ # ls” Mean

The typographical conventions used in this guide are explained in Table 1.1.

Text layout	Meaning
LINUS TORVALDS	important people
GNU Emacs (emacs)	the program GNU Emacs is started with the command emacs
Adabas D	the product Adabas D
/etc/passwd	file or directory name
<file>	when entering a command, <code>file</code> should be replaced by the actual value. Do <i>not</i> enter the angle brackets!
PATH	the environment variable with the name PATH
192.168.1.2	value of a variable
ls	the command ls
'news'	the user 'news'
earth:~ # susehelp	'root' shell in the directory ~ ('~' stands for the “home directory” of the respective user). Enter the command susehelp . Here earth is an example of the name of a computer (host name).
newbie@earth:/tmp > ls	in the shell of the user 'newbie' in directory /tmp, type the command ls
C:\> fdisk	at the DOS prompt, type the command fdisk
Alt	key to press, here the key “Alt”
Ctrl + Alt + Del	when '+' is used, press all listed keys at the same time. Keys to be pressed sequentially are separated by a space only.
"Permission denied"	system messages
'System update'	menu entry 'System update'
Duesentrieb	company “Duesentrieb”

Table 1.1: Typographical conventions used in the text layout

1.4 Acknowledgments

Besides everyone who has contributed to the large success of Linux, we would like to thank EBERHARD MOENKEBERG, who was very helpful with CD-ROM driver problems.

Many thanks go to Hans Lermen, author of the well-known loadlin.exe, which

allows Linux to start from DOS and the DOS `setup.exe` of this Linux system.

Thanks also to all those who sacrificed their spare time for making Linux available to a vast number of international users by helping translating YaST and `linuxrc` texts: Jesús Bravo Álvarez, Gunay Arslan, Jaime Badiella i Aguilera, Zbigniew Baniewski, Sándor Bárány, Olaf Borkner-Delcarlo, Michael Bravo, Görkem Çetin, Franca Delcarlo, Jochen Depner, Benedek Hermann, Ibán José García Castillo, Dora Georgarou (with Romy the barking dog), Jonas-Gintaras Gocentas, Pablo Iranzo Gómez, Krzysztof Hotiuk, Milan Hromada, Arthur Kerkmester, Ralf A. Lanz, Françoise Lermen, Zoltán Levárdy, Nuno Lima, Francisco Javier Moreno Sigüenza, Tomas Morkus, Karine Nguyen, Matts Nordman, Aleksey Novodvorsky, Razvan Cristian Oprea, Gerco Oudhof, Jordi Jaen Pallares, George Papamichelakis, Alexey Pavlov, Ákos Rátkai, Bruno David Rodrigues, Voula Sanida, Gabor Sari, Aleksey Smirnov, Laszlo Varadi, Steve Varadi, P. Vlachodimitropoulos, Joao Teles, Nuno Vieira, January Weiner, Monika Weiß and I Made Wiryana. The meticulous yet riveting task of coordinating this all was undertaken by Antje Faber.

Many thanks, of course, to the dauntless beta testers who risked their running systems: Andreas Koegel, Christian Hüttermann, Dirk Ulbrich, Eberhard Mönkeberg, Frank Hofmann, Georg C. F. Greve, Harald König, Harald Wieland, Karlo Gross, Jens Frank, Lutz Pressler, Martin Hehl, Martin Konold, Martin Schulze, Michael Kleinhenz, Norbert Eicker, Oliver Zendel, Ralf Geschke, Stefan Bliesener, Thomas Wörner, Ulrich Goebel, Ulrich Windl, Volker Lendecke, Wolfgang Barth, and the KDE Team.

Matthias Eckermann has reworked the \LaTeX macros and modernized the layout of the book according to the drafts of Manuela Piotrowski. The new macros are based on the KOMA-`Documents` class by Frank Neukam and Markus Kohm and the “old” `susebuch` macros by Werner Fink and Martin Scherbaum.

The following were involved in the production of the previous manual texts: Bodo Bauer, Daniel Bischof, Michael Burghart, Roland Dyroff, Christian Egle, Klaus Franken, Jürgen Geck, Olaf Hering, Marc Heuse, Florian La Roche, Ralf Lanz, Volker Lendecke, Hans Lermen, Arndt Mehlhorn, Jay Migliaccio, Norbert Munkel, Jordi Jaen Pallares, Christoph-Erdmann Pfeiler, James Samuel, Martin Scherbaum, Burchard Steinbild, Jörg Strelbel, Thorsten Wandersmann, Udo Weber, and Stefan Wintermeyer.

Valuable references — and even partially for the text material: Marco Banse, Stephan Barth, Manfred Blumenthal, Bettina Botter, Georg Brand, Folkmar Claußen, Jörg Dippel, Stefan Dirsch, Felix Domke, Uwe Drechsel, Christian Eggers, Ulrich Felzmann, Kurt Garloff, Lenz Grimmer, Michael Hager, Torsten Hallmann, Michael Hasenstein, Ulrich Hecht, Stefan Hundhammer, Kevin Ivory, Andreas Jaeger, Thomas Jansen, Klaus Kämpf, Bernd Kaindl, Karsten Keil, Wolfgang Kern, Andi Kleen, Sabine Kleine, Thorsten Kukuk, Werner Kurk, Hans Lermen, Marco Michna, Sebastian Nagel, Rüdiger Oertel, Siegfried Olschner, Dirk Pankonin, Petr Pavlík, Rainer Preschke, Andrea Ratschker, Marion Reyzl, Anja Schambach, Christian Steinruecken, Sandra Wiegratz, Steffen Winterfeldt, and Heinz Wohlers.

Holger Dyroff, Christian Egle, Claudio Taulli, and Nicole Zwerenz provided most of the production engineering support. Thanks and keep up the good work!

And the many important people without whom this would never have been possible: to all, a big, big THANK YOU!

*
* *
*

The global team of Linux developers is still working intensively on Linux — most of them on a voluntary basis. We want to thank them for their efforts. This CD distribution would not exist without them. Our work aims at making their work accessible to a large number of interested users.

Last but not least, special thanks to PATRICK VOLKERDING for supporting our work and, of course, many thanks to LINUS TORVALDS!

Last, but not least: Frank Zappa and Pawar!

Nuremberg, 25th October 2001

Have a lot of fun!

Your SuSE Team

2 Your First SuSE Linux Installation

2.1 Installation with the Text-Based YaST (YaST 1)

Use this chapter if, for some reason, you cannot or do not want to use the graphical installation tool, YaST2. For information about installing with YaST2, refer to the *Quick Install Manual*. This chapter also provides in-depth information about partitioning and similar topics, which can supplement the information in the *Quick Install Manual*.

In SuSE Linux, you are provided with YaST, an installation tool which guides you safely through preparing the hard disk, setting up software, and configuring a graphical login.

For the “normal” case of a computer on which SuSE Linux is to be installed, the following requirements must be fulfilled:

- Boot from either CD 1 or from the boot disk that came with your system.
- Sufficient space for the Linux installation must be available on your hard disk. Alternatively, provide Linux with its own sufficiently large hard disk (preferred).
- The CD-ROM drive is supported by Linux. If you are not sure about this, you can find out.

If your hardware does not meet one of these conditions, there are “alternative” methods by which you can successfully complete an installation. These are discussed in the section at the end of this chapter (Section 2.3 page 26).

2.1.1 Starting: The Welcome Screen

Turn on your computer and place CD 1 or the boot disk in the appropriate drive. If the computer refuses to boot, you probably need to change the boot sequence in your system BIOS to either `A, C` or `CD-ROM, C, A`.

After a few seconds, you will be greeted with the start-up screen. You have ten seconds to select ‘Manual Installation’ and prevent an automatic start of YaST2. Add boot parameters to the `boot options` if necessary.

You will know that the loading sequence has started when the following text appears at the bottom of your screen: “Loading initrd...” followed by “Loading linux...”. Then you will see all of the kernel text messages scrolling by and, finally, the program `linuxrc` will start.

The program `linuxrc` is menu-driven and will wait until you enter something.



Figure 2.1: The SuSE Linux Welcome Screen

Possible problems

- If your CD-ROM drive (ATAPI) hangs when the system boots, take a look at Section 2.4.7 page 34.
- If it happens that CD 1, which contains a kernel optimized for Pentium processors, is not detected as a bootable medium, try using the “boot disk” or CD 2. See Sections 2.4.2 page 31 and 2.4.3 page 31.
- With graphics cards like FireGL 1, 2, or 3, the installation is not started in the graphical mode (framebuffer). The installation has to be carried out in text mode.
- As a remedy for most other boot problems, use kernel parameters. See Section 11.3.2 page 239 to learn more. In some cases, the trouble might be DMA related. Try using the ‘Safe Settings’ option offered to circumvent this.

2.1.2 The Basics: linuxrc

The program `linuxrc` allows you to perform settings for the installation. If necessary, load drivers as kernel modules. It will then start the installation program YaST, which, in turn, installs the operating system software and applications.

Tips on operating linuxrc: The \uparrow and \downarrow keys are used to select menu items. Likewise, the \leftarrow and \rightarrow keys are used to select commands, such as 'Ok' or 'Back'. The $\leftarrow\rightarrow$ key will execute the selected command.

A more detailed description of linuxrc can be found in Section 13.5 page 292.

Settings

The program linuxrc now starts by selecting the language, color display, and keyboard.



Figure 2.2: Language Selection

- Select a language. Make sure you highlight 'English' before pressing $\leftarrow\rightarrow$.
- Select either 'Color' or 'Monochrome' (black and white), then press $\leftarrow\rightarrow$.
- Select the keyboard settings. Make sure, once again, that the proper setting is selected before pressing $\leftarrow\rightarrow$.

Adding Hardware with the Modules

You will now be in the linuxrc main menu (Figure 2.3 on the next page) where you have the following choices:

'settings' — Modify the language, screen, or keyboard settings. We already did this.

'System Information' — This menu option shows lots of information about hardware that the kernel has already recognized, modules that have been already loaded, and similar details.




Figure 2.3: The linuxrc Main Menu

'Kernel Modules (Hardware Drivers)' — You might need to enter something here to load modules needed by your hardware.

General rule: You do *not* need to call up this menu item if both your hard disks and CD-ROM drive (ATAPI) are attached to an (E)IDE controller. (E)IDE support is built into the kernel.

'Installation / starting the system' — Continue the installation.

'Abort / Reboot' — Abort the install and reboot your system.

Select loading of kernel modules with `'kernel module'` only if you need support for SCSI or for PCMCIA (see Section 8.1.6 page 190) or if you do *not* have an ATAPI drive; you can read about how modules are loaded in Section 13.5 page 292. In the following submenu choose which modules you want (or need) to load. Possible choices are:

A SCSI module – if you have a SCSI hard disk or SCSI CD-ROM drive.

A CD-ROM module – if your CD-ROM drive is *not* connected to the (E)IDE controller or the SCSI controller.

A network module – if you want to install via NFS or FTP, but this will not be discussed here.

Tip



If you can't find any support for your installation medium (proprietary CD-ROM drive, parallel port CD-ROM drive, network cards, PCMCIA) among the standard modules, you may find what you need on the modules floppy disk which contains additional drivers. To do this, go to the end of the list and select the item `'-- More modules --'`; you are then asked to insert the modules.

Starting the installation




Figure 2.4: Installation Menu of linuxrc

In this menu, you have the following choices:

‘**Start installation**’ Begin installing a new system.

‘**Boot installed system**’ This option is useful if you are having trouble booting to an already installed system.

‘**Start rescue system**’ Up to now, this item is only available for X36-compatible systems.

‘**Eject CD**’ Eject CD.

For the actual installation, press \leftarrow at the menu option ‘Start installation’. In the following screen, select the source medium. By default, the menu option ‘CD-ROM’ is already selected. Now press \leftarrow . The installation environment will be started from CD 1.

As soon as you have completed this step, select the text-based YaST1 (Figure 2.6 page 13) to start the actual installation.

Possible problems

- The SCSI adapter was not recognized:
 - Try the module of a compatible driver.
 - Use a kernel that has built-in support for your SCSI adapter. You need a boot disk as described in Section 2.4.4 page 32.




Figure 2.5: Selecting a Source Medium in linuxrc

- The ATAPI CD-ROM drive hangs while reading information from the CD. In this case, refer to Section 2.4.7 page 34. For more general information, see the section on **General Hardware Problems** in the chapter on **System Administration** in the *Configuration Manual*.
- For reasons which have yet to be clarified, problems can occur when loading data to the RAM disk, which can cause YaST not to be able to be loaded. In such cases, the following remedy usually leads to a positive result:
In the linuxrc main menu select 'Settings' → 'Debug (Experts)'. Set 'Force root image' to no. Return to the main menu and start the installation again.

2.1.3 Installation with YaST

If you selected YaST instead of YaST2 earlier, the YaST welcome screen appears and the reader will see the following four menu options (see Figure 2.7 page 14):

Note

If you want to learn more about YaST or have special maintenance tasks to perform, please refer to the detailed YaST chapter (Chapter 3 page 43). There are hints at the beginning of this chapter on how the keyboard layout works in YaST.



'Install Linux from scratch' If SuSE Linux is to be installed for the first time. This is the topic of the section below.

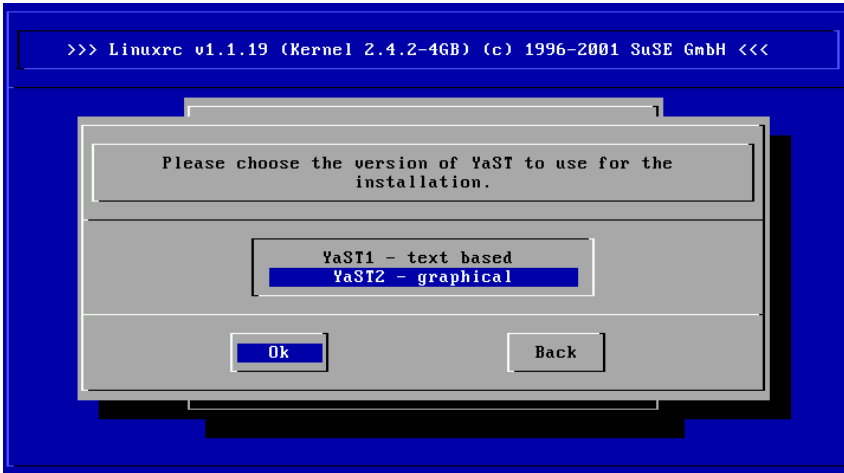


Figure 2.6: Selecting YaST in linuxrc

'Update existing Linux system' Updating a SuSE Linux system is described in Section 12.1 page 271.

'Installation using expert mode' If you select this option, you will have a number of choices to make during the installation. It is strongly recommended that you only choose this option if you are an experienced Linux user and you are sure of the steps necessary for a successful installation. The Expert mode will *not* be covered in the scope of this manual.

'Abort - no installation' if you are having second thoughts...

Select the menu item 'Install Linux from scratch'.

Please take a look *before* you make partitions at Section 2.5 page 36; there you will find useful background information.

2.1.4 Partitioning and Formatting the Hard Disk

YaST now informs you that a hard disk has been detected. If "free" hard disk space is available, YaST will confirm this and suggest that you use this area for Linux (window 'Use Free Area?').



Tip

If you already have a swap partition on your system, perhaps from a previous installation, YaST will ask if this swap partition is to be activated and integrated into the system.

Proceed as follows to set up the partitions:




Figure 2.7: Initial YaST Screen

‘Partitioning’ — As a rule, the question about partitioning for a SuSE Linux first-time installation or a new installation needs to be answered with ‘Partitioning’. You should not select ‘Set up LVM’ unless you know what a *Logical Volume Manager* is and that you require it.

‘Yes’ — If you answer ‘Yes’, YaST will perform the partitioning on its own and you can continue with Section 2.1.6 page 16.

‘No’ — If you say ‘No’, you can partition interactively.

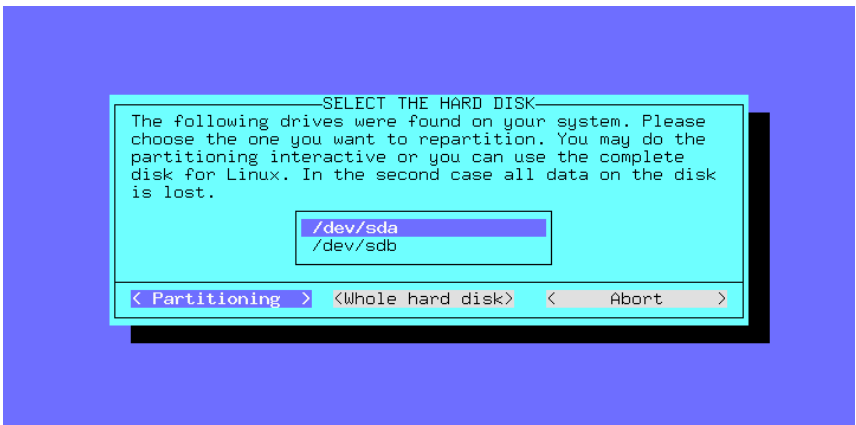


Figure 2.8: YaST — Partitioning Hard Disk


If YaST does *not* find free space on your drive, it will suggest that you ‘Use entire hard disk’. On this screen you can choose from one of the following

possibilities. (A reminder: the `(Tab)`, `(←)` or `(→)` keys can be used to position the cursor, and the `(↵)` key can be used to execute the selection):

'Partitioning' — This lets you partition the hard disk yourself; you will need to select this option if there is another operating system (or partition that you do not want to remove) already on your hard drive.

Caution

If you select 'Use entire hard disk', all data that is currently on the drive will be *lost* and unrecoverable.



'Use entire hard disk' — With this option YaST will — after presenting a big “red” warning screen — *automatically* partition the drive.

You should only choose this option if you just want to install SuSE Linux from scratch, *and* you don't want to get involved with the topic of partitioning.

Operating systems which exist on your hard disk will be deleted if you choose this option!

If you choose the option 'Use entire hard disk', YaST uses the following guidelines when partitioning your hard disk:

- Create a `/boot` partition (a minimum of 2 MB, or 1 cylinder)
- a swap partition. The required size depends on what the machine is to be used for; you will be on the safe side if you can reserve double the amount of RAM and you will not need to worry about any performance limitations. However, if a workstation is well endowed with memory space and you are sure that the applications will not push this to the limit, a smaller value can then be chosen.
- Create a large root partition (`'/'`) with the remaining space.

Possible problems

- It may be the case that there are problems when rereading the partition table. If this is the case, YaST will display this in a *red* display box and ask you to reboot. You should then restart the computer and not attempt to make new partitions the next time.
- Not enough space is available. In section [2.4.1](#) page 28 there are hints on how you can create more space on your hard disk, using `fips` under DOS.

2.1.5 Installation of Software Packages

At this point the hard disk has been prepared for your Linux installation. Now you must decide which software packages to install.




Figure 2.9: YaST — Package Selection

The next YaST menu to appear is the YaST installation menu (Figure 2.9). This may take a few moments to appear because the series and package descriptions must be read from the installation medium and evaluated. If you do not want to choose specific packages, you can continue with the item ‘Start installation’; in this case you can move on to Section 2.1.7 on the facing page.

- Select the menu option titled ‘Load configuration’ to make your software selection (standard, network server, minimal, etc.) This is covered in Section 2.1.6.

When you have selected the system configuration and confirmed this, you will return to the installation menu.

- The menu option ‘Change/Create configuration’ can be used to change “individual” software packages in the list of currently selected packages. You will be presented with the series selection screen. Since you have chosen one of the preconfigured system configurations, it is generally *not* necessary to modify the list of selected packages. The reader should also be aware that it is possible (and easy) to install or remove packages after the system has been installed.

A detailed description of adding packages can be found in Section 3.5.3 page 63.

2.1.6 Selection of the Base Software

The YaST screen for selecting one of the default configurations (Figure 2.10 on the next page) will appear only if you select the menu option ‘Load configuration’, in Section 2.1.5 on the preceding page.

We have prepared a few predefined package “configurations”. With the arrow keys \uparrow and \downarrow you can reach the individual items; with $\text{\textcircled{Space}}$ you can select or

de select an item. An 'x' in the brackets indicates that the configuration has been selected to be added or to replace a currently installed package configuration. A star ('*') next to the configuration title indicates that this configuration is currently selected for installation.




Figure 2.10: YaST — Configuration Selection

'Add' This menu option will add the list of packages from the selected package configuration menu option to the currently installed package configuration.

'Replace' With this option you can replace the currently installed package configuration list with the selected package configuration. You may be asked if you want to delete packages which do not belong to the configuration.

'Abort' This menu option will bring you back to the previous screen.

2.1.7 Installing System Software and Programs

After all of this preparatory work we are finally ready to fill the hard disk with great software!

- You can start installing the software packages with 'Start Installation'. On the screen, YaST will show you the status of each package as it is installed, as well as the total number of packages installed and the number of packages remaining.
- You may be prompted for additional CD's depending on the configuration.
- When the packages have all been installed, return to the main menu by selecting the menu option 'Main menu'.

Possible problems

- If your system has just a “small amount” of RAM (<16 MB), then initially only the packages from the first CD-ROM can be installed. The packages from the other CD’s will be installed later (Section 2.1.9 page 22).
- If you are having trouble copying the software, this is usually caused by hardware problems.

If you have a SCSI system, make sure you double-check the cables and termination. It might also be helpful if you remove all external devices, scanners, etc. from the SCSI bus during the installation. You should also consider using kernel parameters: the most important ones are listed in Section 11.3.1 page 238.

In case you are having problems with an ATAPI device, please refer to Section 2.4.7 page 34.

2.1.8 Selecting a Kernel for the System

The base system has been successfully installed on the hard disk.

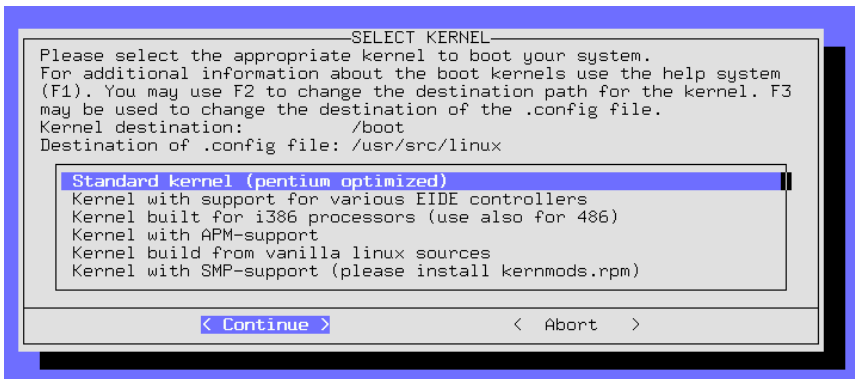


Figure 2.11: YaST — Selecting the Kernel


Now you must install the appropriate kernel. For the system, the boot loader, LILO is set up on the hard disk, and the first software and hardware components will be configured.

- The ‘standard kernel’ is suitable for most systems.
- The ‘kernel for computers with 386 or 486 processors’ can be used for 386 and 486 machines, as well as for machines based on other types of processors not compatible with Pentium class CPU’s.
- The ‘kernel built from original sources’ can be used as a reference which is the main reason why it has been included. Apart from that, you should only select this kernel in cases where the standard kernel shows some kind of erratic behavior.

- The 'kernel with SMP support' should be used for multiprocessor systems.

Caution

If your processor is not 100% Pentium, please select 'kernel for computers with 386 or 486 processors'. You must use this kernel for a Cyrix 686, for example. Selecting the wrong kernel could lead to a "kernel panic"; see also <file:///usr/share/doc/sdb/de/html/cyrix686.html>.



YaST will copy the selected kernel to `/boot/vmlinuz` and the kernel configuration file to `/usr/src/linux/.config`. This file exactly describes the installed kernel and the modules belonging to it.

The question 'Would you like to configure LILO?' should only be answered with 'yes' if you're *positive* that the installed system can be booted with LILO (see Figure 3.26 page 73); generally speaking this is the case for DOS and Windows 95/98 — but for Windows NT things look a little different (see Section 5.7.2 page 115) !

Tips on filling out the necessary forms can be found in Section 3.7.2 page 72.

2.1.9 Base System Configuration with YaST

After selecting the kernel you must finish the rest of the basic configuration. The first software and hardware components are configured.

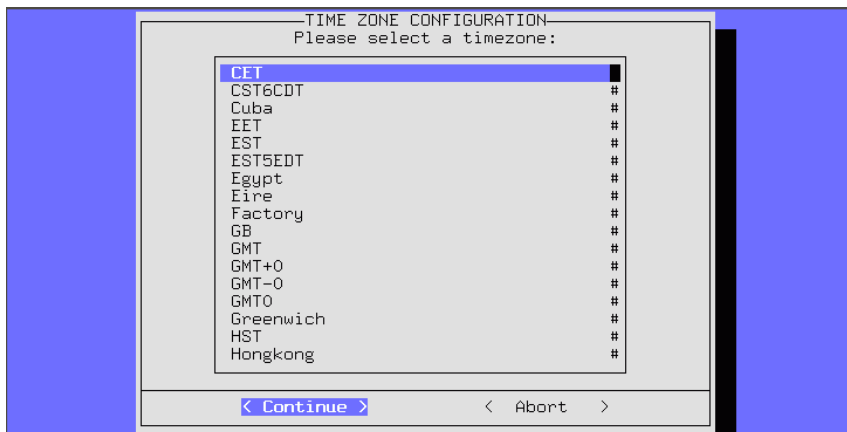


Figure 2.12: YaST — Selecting a Time Zone

- Select the appropriate time zone (Figure 2.12). You will have to search through the long list of time zones to find the one that is appropriate for

your location. You will also be prompted to select either 'GMT' or 'Local time'. You should select 'Local time' unless your system clock (the one in your BIOS) is already set to Greenwich Mean Time.

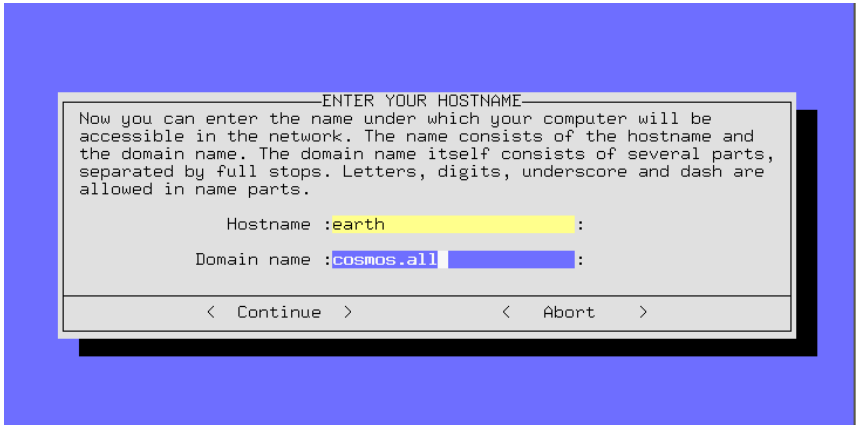


Figure 2.13: YaST — Host and Domain Name

- The following screens that appear are for setting up your network configuration:
 - Host and Domain name (Figure 2.13). Unless you have been given a Fully Qualified Domain Name (from either your System administrator or your ISP), you can choose whatever you like for a Host and Domain name. If you are planning on connecting to the Internet or an intranet, you should make sure that you correctly specify a name. For example earth.cosmos.com. — where earth is the hostname and cosmos.com is the domain name.
 - Onlyloopback, or Real Network? If your computer has *no* network card, you can select loopback and you will *not* be prompted for a detailed network configuration.
 - If you have selected Real Network, you will be prompted with additional network-related questions: DHCP client, Network type (for Ethernet cards, select eth0), IP address, Netmask, Gateway, inetd, portmap, NFS server, From address for Newsgroup postings, Name Server, IP address, selection of kernel module for network card support and Sendmail configuration (sendmail.cf).
- Selecting the sendmail.cf for the mail system (Figure 2.14 on the next page).

The details collected until now are saved and written to the various configuration files by SuSEconfig. YaST terminates and various messages on the state of the system appear on the screen...

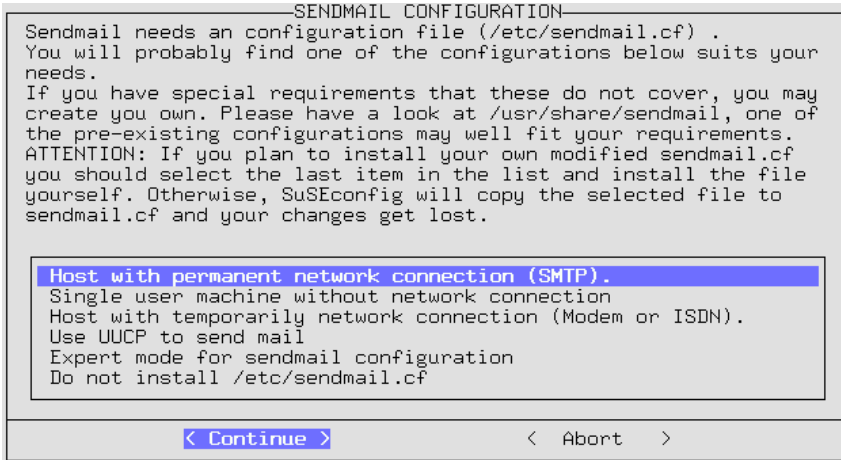


Figure 2.14: YaST — Selecting a Mail System

- After some system messages, you will see a welcome message, "Welcome to SuSE Linux". Here your full attention is required: you are prompted for the 'root' password (System administrator). You should choose this password carefully and not forget it. You should not use "empty spaces" or special characters (unless you know what you are doing). Please note that only the first 8 characters are evaluated.

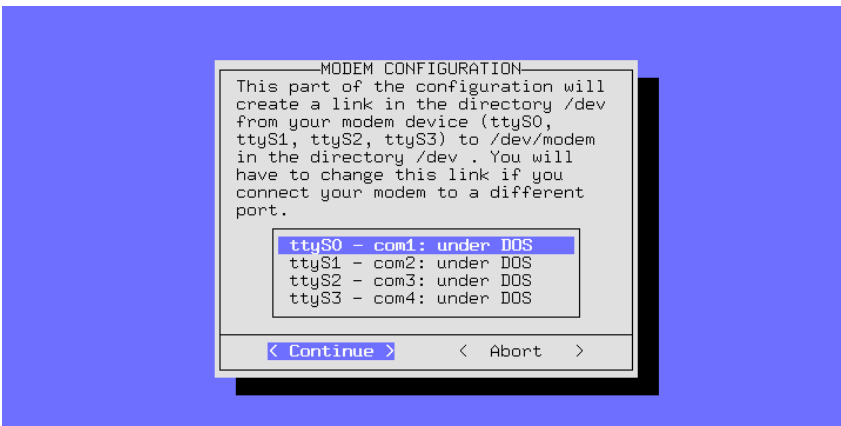


Figure 2.15: YaST — Selecting Interface for Modem and Mouse

- YaST will ask you if you want to configure your modem. If you have a modem, you can do this now or you can configure the modem later. If you choose 'Yes', YaST will prompt you for the serial port to which the modem

is attached. — If you have a “Winmodem” see `file:/usr/share/doc/sdb/de/html/cep_winmodem.html`.

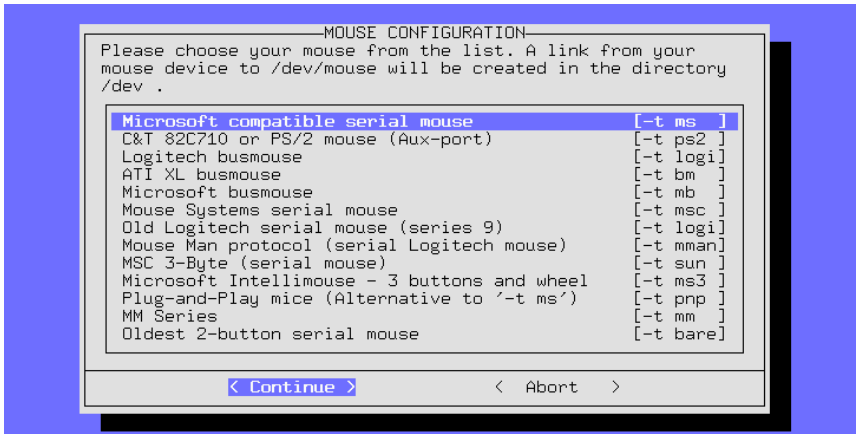


Figure 2.16: YaST — Selecting Mouse Driver

- YaST will then ask you if you want to configure your mouse. If you want to, then select ‘Yes’. Choose your mouse type from the list that appears (Figure 2.16). If you are using a serial mouse, YaST will also prompt you for the appropriate serial port (Figure 2.15 on the page before). Choose the correct serial port from the list.

If additional packages need to be installed from the CD-ROM’s (see Section 2.1.7 page 17), YaST will do this now...

- YaST will ask you to insert the other CD’s, so that the remaining software can be installed; the last CD will only be needed if you have *explicitly* chosen to install the package source code — this is not normally the case!
- YaST will exit now and prompt you to press `(←)`.

The installation of your SuSE Linux is now complete.

2.1.10 Logging in After Your First Installation

Some ‘Root’ configuration scripts will be run in the background. At this point you can log in as ‘root’ in the foreground, for example on Console 1. **Login:** At the login prompt “Login:”, enter `root` and you will be prompted for the “Password:”. Here you should enter the password that you gave for the ‘root’ user (see Section 2.1.9 on the page before).

The Linux prompt appears and you can now start working; this is what the prompt looks like:

```
earth:~ #
```

For example, the command `ls -a` can be used to view the contents of the current directory:

```
earth:~ # ls -a
```

The program `yast` will start YaST, which can be used to make changes to your system configuration:

```
earth:~ # yast
```

If you start YaST (`yast`) you can select the menu option ‘System Administration’ and the submenu ‘User administration’ to create new user accounts. You can continue with the subitem ‘Configuring XFree86™’; via this item the graphical interface can be set up (see also Section 6.3 page 132).


Configuration scripts will be running automatically in the background (indexing man pages, setting up Perl, etc..). On computers with limited memory and slow CPUs, this procedure can take up to an hour to complete. If you decided to “reboot” your computer before these scripts are finished, YaST will automatically continue processing the scripts from where it left off!

You can view the status of these scripts on Console 9. To see this console, press the `(Alt) + (F9)` keys simultaneously. Once the scripts are completed, the following message will appear:

"Have a lot of fun!"

After the installation is complete, additional configurations can be made.

2.1.11 Shutting Down and Booting the System



Caution

You should never switch off a Linux computer directly. Either use the function provided by KDM or the commands presented in Section 2.1.11 to “shut down” the computer. You can switch it off after shutdown. If the computer is not shut down properly, a file system check must be carried out the next time the machine boots — this takes some time and in exceptional cases, the check is not sufficient to restore all files which may have been damaged.

The commands for shutting down your system properly are shown in Table 2.1.

```
shutdown -h now halts the system; when you see the output: "the
system is halted", you can switch off your ma-
chine
```

```
shutdown -r now reboots the system immediately
```

Table 2.1: Commands for shutting down your Linux system

The **shutdown** command can only be invoked by `'root'`.

To shut down your system, log in as `'root'` and enter the commands **shutdown -h now** or **shutdown -r now**.

Alternatively, you can use `(Ctrl) + (Alt) + (Del)`, which you might already know from booting DOS. This combination does not work if you are currently running X11. Nevertheless, this method can be used by any user on any virtual console.

2.2 Starting SuSE Linux

After installation, the question remains as to how you want to boot your system normally (booting).

The following section describes the options you have for booting your SuSE Linux system. The best boot method for you will depend on the intended use of your Linux system.

Boot disk: You can boot Linux from a *boot disk* (“boot floppy”). This choice will always work (as long as the boot disk does not get corrupted) and it is easy. The boot disk may have been created during the installation.

loadlin: The loadlin boot option can be configured as follows:

- The computer must be running either in DOS Real mode or have a VCPI server in virtual 8086 mode¹. In other words, this method will *not* function in Unix, OS/2, Windows NT, or in a DOS window on a Windows 95/98 machine. It does, however, function well from MSDOS or from Windows 95/98 in DOS mode.
- Your computer must have enough DOS memory available: there must be 128 KB available below the 640 KB limit, the rest can be on extended/EMS/XMS memory.


loadlin is fairly labor intensive to set up, but it can be easily integrated into the Windows 95/98 boot menus. This requires a manual editing of configuration files. One of the biggest advantages of loadlin is that nothing gets installed into the MBR (*Master Boot Record*) of the hard disk. To other operating systems, Linux partitions will appear as an unknown type.

To install loadlin, you need to know a little about Linux and DOS. You should also be able to create configuration files with an editor. Find details in Section 5.9 page 122. If you make a mistake in the Windows 95/98 boot menu, this could cause problems. In the event of an extreme error, you may lose access to your Windows hard drive. *Before* you start modifying your Windows boot menu, make sure that you can boot your Windows operating system with a boot floppy.

LILO: The universal and technically elegant solution for booting your system is LILO. The LILO boot menu can be configured to give a choice of operating systems to boot before anything is booted. Set up LILO during the installation

¹A VCPI server is accessed by the program `emm386.exe`.

or choose to configure it later on with the help of YaST (see Section 3.7.2 page 72). LILO must be installed in the boot sector of your hard disk, which is not without some risk. You will need a basic understanding of how the boot process actually works if any problems arise. You should also be comfortable enough with the main LILO configuration file to edit it. Finally, you are well advised to learn how to uninstall LILO in the event that you have difficulties. Nevertheless, LILO is considered the best boot method available.



Caution

There are BIOS variants which check the structure of the boot sector (MBR), and after a LILO installation erroneously display a virus warning. This problem can be easily removed by entering the BIOS and looking for corresponding adjustable settings. For example, you should switch off 'virus protection'. You can switch this option back on again later. It is unnecessary, however, if Linux is the only operating system you are using.

A detailed discussion of various boot methods, especially of LILO and loadlin, can be found in Chapter 5 page 99.

The Graphical SuSE Screen

Beginning from SuSE Linux 7.2, the graphical SuSE screen is enabled on the first console, as long as the option "vga=771" has been provided as a kernel parameter. If you install using YaST2, this option is added to `/etc/lilo.conf`.

Disabling the SuSE Screen

Basically there are three ways to achieve this:

- Disabling SuSE screen on an as-needed basis:

Enter the command

```
earth: # echo 0 >/proc/splash
```

to disable the graphical screen. By issuing the command

```
earth: # echo 0x0f01 >/proc/splash
```

you can enable the screen again.

- Disabling the SuSE screen by default:

Add the kernel parameter `splash=0` to your boot loader configuration. In most cases, this will mean that the file `/etc/lilo.conf` needs to be edited to include the following line:

```
append="splash=0"
```

If your `lilo.conf` already has an `append` line included in it, just edit this line, for example, by changing `append="vga=771"` into `append="vga=771 splash=0"`. However, if you prefer the old text mode anyway (which was the default with previous versions), you should set this to `"vga=normal"`.

- Completely disabling the SuSE screen:

Compile a new kernel after disabling the option `Use splash screen instead of boot logo` in the menu `'framebuffer support'`.



Tip

Removing the framebuffer support from the kernel will automatically disable the splash screen as well. SuSE cannot provide any support for your system if you run it with a custom kernel.

For most users, the recommended method is to edit `/etc/lilo.conf` as described above. Execute `lilo` afterwards so that LILO reads in the changes in its configuration. Also see Section 5.5 page 110.

2.3 Special Installations

2.3.1 Installation Without a Supported CD-ROM Drive

What do you do if a standard installation via the CD-ROM drive is not possible? Your CD-ROM drive might not be supported if it is an older, “proprietary” model. It might be your second computer, a notebook, for example, which might not even have a CD-ROM drive, but instead has an Ethernet adapter or a PLIP cable. SuSE Linux provides ways of installing over a network connection to such a machine without using a supported CD-ROM drive. NFS or FTP via Ethernet or via PLIP (Section 2.3.1) is available.

Installation from a Source in the “Net”

We do not offer support for this method of installation. It is only recommended for experienced computer users.

What's it all about?

There is *no* CD-ROM installed on the machine on which you want to install Linux. You are able to connect to a remote machine over the network that has an installed CD-ROM drive (and CD) or which has a copy of the CD on its hard disk.

In addition, it is necessary to copy the files `.S.U.S.E-disk*` from the CD-ROM to the hard disk. This can be abbreviated in the following way:

```
earth: # cp /media/cdrom/.S* /emil
earth: # cp -a /media/cdrom/suse /emil
```

This “other” computer must “export” the directory in a suitable manner.



Tip

Alternatively, it is sufficient to start the rescue system on the remote machine and directly export the CD 1.

Step by step

1. Start the installation of the client as specified in Section 2.1 page 7.
2. Continue with the installation as described in Section 2.1.2 page 8, but:
 - when you come to ‘Kernel modules’, select ‘Networking cards’ and load the necessary driver. This is not necessary if you are installing via PLIP.
 - when prompted by `linuxrc` to specify the ‘Source medium’, select ‘Network (NFS)’ then configure the network using the available menus. As an alternative, install via FTP.
3. Finish up the installation as given in Section 2.1.3 page 12.

Possible problems

- The installation aborts before it has actually started, because the installation directory of the “other” machine was not exported with `exec` permissions. Correct this and start again.
- The server does not know the computer on which SuSE Linux should be installed. Enter the name and IP address of the computer which is to be newly installed into the file `/etc/hosts` of the server.

2.3.2 Automatic Installation and Configuration of SuSE Linux

Why Automatization?

The automatic installation and configuration of Linux systems allows you to assemble a unified server landscape. Even for client installations, automatic installation is preferable — above a certain number — to a manual installation. This standardization involves the state of the system versions and the software, file system structures, and configuration files. The automatization guarantees that an installation method which has once proven successful can be used on a computer at any time, without expert knowledge. In this way, it is also possible to expand the server landscape without problems. This standardization is an obvious help for system administrators: if the same configuration structures are required, you do not need to search for configuration files on every server each time an installation is performed. The software behaves as expected, because the same versions are being used. Bug fixes will take effect across the whole server landscape. All this improve the quality of production.

An overview of ALICE

ALICE, *Automatic Linux Installation and Configuration Environment*, integrates the installation and essential areas of configuration. Due to the CVS-based behavior of the configuration data, you can maintain this configuration data from various locations using different users.

It is possible to classify a computer into one or several different classes. This means that different computers can be kept as “similar” as possible. Only really necessary data needs to be given in the computer-specific configuration file. To organize its structure, the configuration data is divided into three different directories. The `templates` directory contains a number of defaults, the `classes` directory contains class-specific configurations, and the `info` directory contains host-specific information. The defaults can be overwritten by class values, and these, in turn, by special values assigned to the host.

ALICE is a growing project. If you do not find any modules for configuring software you are using, this can still be accomplished independently or in the framework of a project.

Further Information

Information on ALICE modules, setting up an installation server, the architecture of ALICE, and a configuration session can be found in the documentation included in the package `alice`.

Automatic installation with YaST1

The “Automatic Installation” of YaST1 can also be used independently of the complete framework of the ALICE project. See also the hints in the article file:/usr/share/doc/sdb/en/html/cg_autoinstall.html.

2.4 Tips and Tricks

2.4.1 Creating Space for Linux (Partitioning)

What’s it all about?


You need to prepare your hard disk for Linux partitions. Hopefully you have some time to read through this section carefully. It is recommended that you have some disks or tapes for a backup and a boot disk to boot your existing operating system.

Details

By partitioning, your hard disk can be separated into several independent parts. One reason for partitioning your hard disk might be the coexistence of several different operating systems with different file systems on your hard drive. Hard disks are divided into cylinders. Each cylinder always contains exactly the same amount of tracks, cylinder 0 being the innermost.

How do you get space for your Linux partition?

- Individual partitions on the hard disk can be deleted, but all data residing on these partitions will be lost. This space is now available for new partitions (e. g., Linux). You could, however, split this partition into two or more partitions.
- Any partition can be assigned a different file system. As with deleting, all data on these partitions will be lost.
- In MS-DOS or Windows 95/98, reduce the last partition of the hard disk in size without losing files in it. Ensure, with a defragmentation program, that all files are located at the beginning of the partition. If you just have *one* MS-DOS or Windows partition, you can make space in this way for Linux partitions very simply. After defragmenting, you can, with special programs, for example, the Freeware program, *fips*, reduce the cylinder upper limit, making the partition smaller. Find *fips* on CD 1 in the directory *dosutils*.
- The easiest way is to buy a new hard disk. This, of course, will cost some money.



Caution

Every change of the partition tables must be done with extreme caution. Read the documentation of the corresponding programs. Severe problems can occur while changing partition tables. You can even lose all your data. SuSE cannot be held responsible for this. It is strongly recommended that you have at least a boot disk and a backup of your important data.

Step by step

To partition your hard disk:

1. If you do not know already, find out how many partitions are on your hard disk and the size of the partitions. Use *fdisk* of your OS for this task.
2. Determine how many partitions you really need and which size you are going to assign them. Information on this can be found in Section 2.5 page 36 and in Section 2.6 page 39.
3. Write down the partition data. You need it later during the installation process.
4. A backup of your hard disk at this point is highly recommended. If you do not have a backup tape or streamer installed and do not want to back up everything onto floppies (which is a tedious job), at least make a backup of the most important data and files (e. g., *autoexec.bat*, *config.sys*, **.ini*, etc.). Create a boot disk for your OS **and** make sure you can boot with it. There are several programs that might come in handy on your boot disk, such as *fdisk*, an editor, a formatting program, and a backup program.
5. The next step depends on your system configuration.

DOS/Windows 95/98, one partition on your hard disk, and no backup

You have to reduce the size of your partition without deleting data. Move all data to the beginning of your partition using a tool such as defrag (MS-DOS 6 or Windows 95/98).

The defragmentation program does not normally move hidden or system files because sometimes they are write-protected and have a defined place on the hard disk. Moving them could have adverse effects. If you are sure that no such data is on your hard disk, deactivate the attributes `hidden` or `system` for the necessary files or, depending on which defragmentation program you are using, “force” the defragmentation of such files. Keep in mind that even the Windows swap file is a hidden file. If it is in the way, deselect it in Windows with “disable virtual memory”.

Once you have retrieved enough space, change to the directory `\dosutils\fips15` on your SuSE CD. There you will find `fips.exe`, with which you can make partitions smaller. Study the documentation included with `fips` carefully since changing partition tables can be quite tricky. The program `fips.exe` only runs on DOS, not on Windows. Exit Windows or, if using Windows 95/98, change to DOS mode. After running `fips.exe`, you will have a second partition which later will be divided into your Linux partitions.

Tip

`fips.exe` in `\dosutils\fips\fips20` is capable of shrinking `fat32` partitions. Make a backup *before* trying out this `fips` version.



DOS/Windows 95/98 with several partitions or you have a complete file backup

Delete your DOS partitions and, afterwards, install them again, reduced in size. You are going to lose all data on those partitions, so you have to have a complete backup available (file backup, not image backup). Using `fdisk`, delete all partitions and install the new ones. Next, format these partitions, install your OS, and copy all data from your backup back to disk (remember that you must have your backup program on your boot disk).

OS/2

Here, you have the following possibilities:

- Reduce your OS/2 partitions. To accomplish this, you must backup all data on your OS/2 partition, delete these partitions using OS/2's `fdisk`, and reinstall them reduced in size. The space now available will be used later for your Linux partitions.
- If you plan to use Linux in future instead of OS/2, change the partition type later to Linux or delete all OS/2 partitions and recreate them later as Linux partitions.

- Buy a new hard disk. Formatting and partitioning it will be done later with YaST.
6. Reboot your computer.
 7. Check whether your old system is still running. Linux partitions on the freshly created new space will be installed later using YaST.

Possible problems

The following problems can occur:

- The partition might not be adequately reduced in size because defrag cannot move some files.

2.4.2 Booting from Disk (SYSLINUX)

The “boot disk” method can be used as a fallback solution if you have to deal with installation requirements which are somewhat special (e. g., the CD-ROM drive is not available).

The bootstrapping process is initiated by the bootloader SYSLINUX (package `syslinux`). SYSLINUX is configured so that some minimal hardware detection is performed on bootup. Basically this includes the following steps:

- Checking whether the BIOS provides VESA 2.0 compliant framebuffer support and making sure the kernel is booted accordingly.
- Reading out any available monitor information (DDC info).
- Reading the first sector of the first hard disk (the “MBR”) to arrange for the LILO configuration later, where BIOS IDs will be assigned to Linux device names. This involves a test to read out the sector using the BIOS’s `lba32` function, to find out whether the BIOS supports the corresponding features.



Tip

If you keep `(Shift)` pressed when SYSLINUX is started, all of the above steps will be skipped. To ease troubleshooting, add the line `|verbose 1|` to your `syslinux.cfg`, in which case the bootloader will tell you which action is being performed.

2.4.3 Booting with CD 2

As well as CD 1, the second CD is also bootable. On CD 1, a 2.88-MB boot image is used. The second CD contains a traditional image of 1.44 MB in size. Use CD 2 when you know for certain that you can boot from CD, but when things do not work with CD 1 (the “fallback” solution). Unfortunately, not every BIOS correctly recognizes the large images.

2.4.4 Creating a Boot Disk In DOS

Requirements

You need a formatted 3.5" floppy disk and a bootable 3.5" floppy drive. If you are working in Windows 95/98, launch `setup` from MS-DOS mode, *not* from inside a DOS window.

Additional information

CD 1, in the directory `/disks`, contains a number of disk images. Such an image can be copied to a disk with the help of suitable auxiliary programs. This disk is then called a boot disk. Included in these disk images are the “loader”, `SYSLINUX`, as well as the program `linuxrc`. `SYSLINUX` allows selection of a specific kernel for the booting process and to add parameters for your hardware, if necessary. The program `linuxrc` supports the loading of kernel modules for your hardware then starts the installation.

Normally the SuSE boot disk supplied can be used to boot. Only for exotic hardware not supported by the modularized kernel of this boot disk or if you download a disk image from the Internet (for example, from <ftp://ftp.suse.com>), do you need to create your own boot disk as described here.

With Setup

Step by step

To create a boot disk:

1. Start `setup` directly from CD 1.
2. Select ‘`floppy`’ and press `(←)`. Next, select ‘`Boot`’ and press `(←)`.
3. Select a disk with a suitable kernel, for example, that supports your SCSI adapter if you have one. `setup` shows the essential part of the kernel descriptions. If you need further information, look it up in `\disks\readme.dos`. Remember the name of your kernel. You will need it later. Now press `(←)`.
4. Create the boot disk. Insert the (DOS-formatted) disk into the 3.5" drive and select the disk to create.
 - Only the boot disk is needed (‘`Root`’ is not needed anymore for SuSE Linux.) Move the cursor onto ‘`Boot`’ and press `(←)`.
 - `setup` requests confirmation of disk insertion. Press `(←)` and the disk is written.
 - When this is finished, press `(←)`.
 - Now, select ‘`Done`’ to exit this screen and `setup`.

With rawrite

Alternatively, you might want to use the (perhaps slower) DOS program `rawrite.exe` (CD 1, directory `\dosutils\rawrite`) to write the disk at the boot prompt.

The standard disk images are contained on CD 1 in the directory `/disks`. Read the file `README`. The image `bootdisk` or `scsi01` is the usual choice for the standard disk. All the actual kernels can be found in the directory `/suse/images` (without extensions). Also read the `README` file there.

If you need the standard disk which is supplied with every SuSE Linux, proceed as follows. It is assumed that you are in the directory of the CD.

```
Q:> dosutils\rawrite\rawrite
disks\bootdisk
a:
```

If you need a specific type of support, another disk image should be used instead of `bootdisk`. If problems arise, `k_i386` can be implemented as a fallback kernel.

2.4.5 Creating a Boot Disk with UNIX

Requirements

You need access to a Unix or Linux system with an accessible CD-ROM drive and a formatted disk.

To create a boot disk:

1. If you need to format the disks first:

```
earth: # fdformat /dev/fd0u1440
```

2. Mount the first CD (disk 1) (e. g., to `/cdrom`):

```
earth: # mount -tiso9660 /dev/cdrom /cdrom
```

3. Change to the `disks` directory on CD:

```
earth: # cd /cdrom/disks
```

4. Create the boot disk with

```
earth: # dd if=/cdrom/disks/bootdisk of=/dev/fd0 bs=8k
```

In the `README` file in the directory `disks`, read about what features specific kernels have. These files can be read with `more` or `less`.

5. If you need a different kernel, another disk image can be used in place of `bootdisk`. If problems arise, `k_i386` can be implemented as a fallback kernel.

Somewhat more complex is using a self-compiled kernel during installation. In this case, write the default image (`bootdisk`) to the disk then overwrite the actual kernel (`linux`) with your own kernel (refer to Section 10.5 page 233):

```
earth: # dd if=/cdrom/disks/scsi01 of=/dev/fd0 bs=8k
earth: # mount -t msdos /dev/fd0 /mnt
earth: # cp /cdrom/suse/images/eide.ikr /mnt/linux
earth: # umount /mnt
```

2.4.6 Does Linux Support my CD-ROM?

Almost every CD-ROM drive is now supported by Linux.

- Using ATAPI drives (those drives connected to an EIDE controller), there should be no problems at all.
- Using SCSI CD-ROM drives, it is only important whether the SCSI controller is supported by Linux (supported SCSI controllers are listed in the Appendix). If your SCSI controller is not supported and there is a HD connected to it, you will have a problem anyway.
- Many proprietary CD-ROM drives are supported under Linux (see Appendix). Here is the most likely place for problems to occur. If your drive is not mentioned, try choosing a similar type.
- In the meantime, parallel port CD-ROM drives have become very popular. Unfortunately, there is no standard, leading to unexpected trouble. SuSE Linux contains a number of alpha drivers for some devices. If none of them works, install via DOS. Keep in mind that you may access some of these devices only after they have been set up under DOS. You might need a warm reboot.

2.4.7 ATAPI CD-ROM Hangs While Reading

If your ATAPI CD-ROM is not recognized or it hangs while reading, this is most frequently due to incorrectly installed hardware. All devices must be connected to the EIDE controller in the correct order. The first device is master on the first controller. The second device is slave on the first controller. The third device should be master on the second controller, and so forth.

It often occurs that there is only a CD-ROM besides the first device. The CD-ROM drive is sometimes connected as master to the second controller (secondary IDE controller). This is wrong and can cause Linux not to know what to do with this “gap”. Try to fix this by passing the appropriate parameter to the kernel (`hdc=cdrom`) (see Section 11.3.2 page 239).

Sometimes one of the devices is just “mis-jumpered”. This means it is jumpered as slave, but is connected as master, or vice versa. When in doubt, check your hardware settings and correct them where necessary.

In addition, there is a series of faulty EIDE chipsets, most of which have now been identified. There is a special kernel to handle such cases. See the `README` in `/disks` of the installation CD-ROM. The kernel parameters needed are described in detail in Section 11.3.2 page 239 and Chapter 10 page 229.



Note

If booting does not work straightaway, try using the following kernel parameters. At the boot prompt (**boot:**), enter:
boot: `linux <relevant parameters>`
 (see below) *Note:* Do not forget to add the name of the kernel (`linux`) before the actual parameters!

hd<x>=cdrom — `<x>` stands for a, b, c, d, etc., and is interpreted as follows:

- a — Master on the first IDE controller
- b — Slave on the first IDE controller
- c — Master on the second IDE controller
- ...

An example of `<parameter to be entered>`: `hdb=cdrom`

With this parameter, specify the CD-ROM drive to the kernel, if it cannot find it itself and if you have an ATAPI CD-ROM drive.

ide<x>=noautotune — `<x>` stands for 0, 1, 2, 3, etc., and is interpreted as follows:

- 0 — First IDE controller
- 1 — Second IDE controller
- ...

An example of `<parameter to be entered>`: `ide0=noautotune`

This parameter is often useful for (E)IDE hard disks.



Note

Further kernel parameters can be found in [Section 11.3.2](#) page 239. In case of problems with SCSI systems or with connecting network cards, you should also look here.

2.4.8 Problems with CD-ROM Drives on Parallel Port

All available drivers are listed by `linuxrc` at the installation. Normally there are no peculiarities. Unfortunately, lots of drives (e. g., Freecom) are not supported yet. It may be that you cannot use them although the manual claims that the type is identical. The manufacturer apparently has changed the internals without making these changes public.

Some devices need to be initialized by the DOS driver before making them available under Linux:

1. Boot DOS and load the CD-ROM driver.
2. Insert a Linux boot disk.
3. Warm-reboot the machine.

If your drive is not supported, as you did before, install from a DOS partition (see Section 2.3 page 26).

For current information on parallel port programming have a look at: <http://www.torque.net/linux-pp.html>.

2.4.9 Loadlin Does Not Have Enough Memory to Load the Kernel

You do not have enough free memory below 640 KB. Try to remove drivers from your startup files or shift them to high memory.

If you use compressed drives under Windows 95/98 and shifting the driver to high memory does not work, decompress those drives.

2.4.10 Loadlin Does Not Start

If you encounter any problems using loadlin, start `loadlin` using the following options: `-v`, `-t`, or `-d`.

It is best to write the debug information into a file, `debug.out`.

```
C:\> loadlin -d debug.out <other parameters>
```

You could send this file to the SuSE support. For `<other parameters>`, enter your system-specific values (see Section 5.9.1 page 123).

2.5 Partitioning Tips

Is Linux and its file systems completely new to you? Then you probably ask yourself questions like: How much space should I assign to Linux? What's the minimum I need? What's best for my needs? How should I divide up the available space?

Partition Types on a PC

Every hard disk contains a partition table that has space for up to four entries. Each entry may be either a primary or an extended partition. Only *one* extended partition may be assigned.

Primary partitions are a continuous section of cylinders which are assigned to one operating system. Using primary partitions, you would only be able to create four partitions. More will not fit into the partition table.

An extended partition itself is also a continuous section of cylinders, but you can divide an extended partition into multiple *logical partitions* which do not need a

separate entry in the partition table. The extended partition is like a container for logical partitions.

If you need more than four partitions, you have to make one an extended partition in which you may assign logical partitions. The maximum is 15 partitions for SCSI systems and 63 for (E)IDE systems.

Linux does not care on what kind of partition it is installed. It can be primary or logical.

Time to Decide

Let's start with the minimum SuSE Linux install: 180 MB. This only works if you use the machine for a simple purpose, e. g., you only work on a text console (no X Window System). If you want to take a look at X and start a few applications, you will need 500 MB. Both values include swap.

What is a reasonable size for an installation? 1 GB — in this world of several gigabyte hard disks, this is quite a modest requirement. There is no upper limit.

What's the best for your needs? That depends on what you want to do:

- Working in X and using applications such as **Applixware** and Netscape will require 1.2 GB.
- To create your own applications in X, you also need 1.2 GB.
- Both of the above items: 2 GB.
- To compile your own X servers and write your own CDs together with the items mentioned above: 4 GB.
- To set up an Internet or FTP server: 700 MB minimum.



Tip

With the robustness of the Linux file system now achieved, it is a very *good* idea, especially for novices, to follow the strategy suggested by YaST: a small partition at the beginning of the hard disk for `/boot` (at least 2 MB, for large drives, 1 cylinder), a partition for swap (64–128 MB), all the rest for `/`.

If you want to partition something, but also as little as possible, then the following simple rules apply.

How should you divide up the hard disk? This simple question cannot be answered easily. Apply these guidelines:

- Up to 500 MB: swap partition and a root (`/`)
- Approx. 500 MB to 4 GB: small boot partition for the kernel and LILO at the *very beginning* of the hard disk (`/boot`, approximately 8–16 MB or 1 cylinder), a swap partition, and the rest for the root partition `/`.

- More than 4 GB: boot (`/boot`), swap, root (250 MB), home (`/home` with approximately 200 MB per user) and the rest for applications (`/usr`); possibly a further partition for `/opt` (see on this page) and for `/var`.



Note

If you plan to start Linux directly from the hard disk, you will need a Linux partition below the “1024–cylinder limit” as a boot partition (read more on this on page 103). This does not concern you if you are starting Linux from DOS/Windows with `loadlin`. Usually, the boot partition (`/boot`) will be the same as the root partition.

Bear in mind that some programs (mostly commercial programs) install their data in `/opt`. If you are installing a number of these, either create a separate partition for `/opt` or enlarge the dimensions of the root partition accordingly. Specifically, this involves the program packages or demo programs listed in Table 2.2, calculated with size increases in mind (in the table mentioned there are also programs which are *not* included in SuSE Linux).

KDE	200 MB
GNOME	200 MB
htdig	5 MB
Fortify	2 MB
susehelp with htdig full text search	40 to 300 MB
Wabi	10 MB
Mozilla	55 MB
Netscape	45 MB
Arcad	350 MB
Applixware	400 MB
Eagle	18 MB
StarOffice	260 MB
Cyberscheduler Software	30 MB
Cygnus Source-Navigator	20 MB
SNiFF+	45 MB
Insure++	45 MB
pep	18 MB
Oracle 8.1.7	900 MB
Sybase — Adaptive Server Enterprise	170 MB
virtuoso — OpenLink Virtuoso Lite Edition	55 MB

Table 2.2: Packages in Directory `/opt`

2.6 Partitioning for Experts

In Section 2.5 page 36 and Section 3.4.11 page 56, details are given of the various ways to partition your system. This section should provide more detailed information for tailoring a system that best suits your needs. This section is mainly of interest for those who want an optimized system as far as security and performance are concerned and who are prepared to reinstall the complete system where necessary.

It is absolutely essential that you have extensive knowledge of the functions of a UNIX file system. You should be familiar with the topics mount point and physical, extended, and logical partitions.

It should be mentioned here that there is *no* golden rule for all, but many rules for each situation. Don't worry, you will find concrete figures in this section to help you.

First, you need to gather the following information:

- What is the purpose of the machine (file server, compute server, stand-alone machine)?
- How many people are going to work with this machine (simultaneous logins)?
- How many hard disks are installed? How big are they and which kind (EIDE, SCSI, or even RAID controllers)?

2.6.1 Size of Swap Partition

Quite often you will read: “Swap should be at least as large as physical RAM”. This is a relic of times when 8 MB was regarded as a lot of RAM Memory.

Applications that need considerable memory have shifted these values up. Generally, 64 MB of virtual swap should be sufficient. Do not be stingy. If you compile a kernel in X and want to have a look at the manual pages using Netscape and have an emacs running, you will already take up all 64 MB.

To be safe, opt for at least 128 MB of virtual memory. One thing you should never do is not assign swap space at all. Even on a machine with 256 MB RAM, there should be a swap partition.

Do you plan to run extensive simulations and need gigabytes of memory? When in doubt as to whether Linux suits your needs, read Section 2.6.2 on the next page (Example: compute server).

2.6.2 Computer Used as Stand-Alone Machine

The most common use for a Linux machine is as a stand-alone computer. To make decisions as easy as possible for you, we provide you with some concrete figures which you can use at home or at your company. Table 2.3 on the following page is an overview of size requirements for different Linux systems.

Installation	Disk space needed
minimum	180 MB to 400 MB
small	400 MB to 800 MB
medium	800 MB to 4 GB
large	4 GB to 8 GB

Table 2.3: Examples of Disk Space Requirements for Different Installations

Example: stand-alone machine (small)

You have a 500 MB spare hard disk to hold Linux: use a 64 MB swap partition and the rest for / (root partition).

Example: stand-alone machine (average)

You have 1,2 GB available for Linux. A small boot partition `/boot` (5–10 MB or 1 cylinder), 128 MB for Swap, 800 MB for /, and the rest for a separate `/home` partition.

Example: stand-alone machine (deluxe)

If you have more than 1.2 GB available, there is no standard way to partition. Read Section 2.6.3 on the next page.

Using as a file server

Here, hard disk performance is *really* crucial. You should use SCSI devices if possible. Keep in mind the performance of the disk and the controller. A file server is used to save data centrally. This data might be home directories, a database, or other archives. The advantage of this is that administration of the data is simple.

If the file server will serve a huge net (20 users and more), optimizing hard disk access is essential. Suppose you want to provide a file server for 25 users (their home directories). If the average user requires 100–150 MB for personal space, a 4 GB disk mounted under `home` will probably do.

If there are 50 such users, you will need an 8 GB disk. In this case, it would be better to split `home` into two 4 GB disks, as now they would then share the load (and access time).

Using as a compute server

A compute server is generally a powerful machine that carries out extensive calculations over the net. Normally, such a machine is equipped with extensive main memory (512 RAM or greater). The only point where fast disks are needed is for the swap space. The same rule applies here: if you have a number of hard disks, you can spread swap partitions across them.

2.6.3 Optimizations

The disks are normally the limiting factor. To avoid this bottleneck, there are two possibilities which should be used together:

- separate the load onto multiple disks
- use an optimized file system (e.g., `reiserfs`).
- equip your file server with enough memory (at least 256 MB)

Parallelizing multiple disks

This needs some further discussion. The total amount of time needed for transferring data can be separated into five factors:

- time elapsed until the request reaches the controller
- time elapsed until this request is send to the disk
- time elapsed until the hard disk manages to set its head
- time elapsed until the media has turned to the right sector
- time elapsed for transferring data

The first factor depends on the network connection and has to be regulated elsewhere. We do not cover this here. The second factor can be ignored, because it depends on the controller. The third factor is the vital part. The time is counted in milliseconds. Relative to the access time of main memory (measured in nanoseconds), this is a factor of one million. The fourth factor depends on the disk rotation speed. The fifth factor depends on the rotation speed, the number of heads and the actual position of the data (inside or outside).

For optimized performance, consider factor three. Here, the SCSI feature *disconnect* comes into play. Look at what happens:

The controller sends the command (in this case to the hard disk) “Go to track *x*, sector *y*” to the device. Now the disk motor has to start up. If this is an intelligent disk (if it supports *disconnect*) and the driver itself is also able to do *disconnect*, the controller sends a *disconnect* and the disk separates itself from the SCSI bus. Now other SCSI devices can do work. After a time (depending on the strategy or load on the SCSI bus), a connection to the disk is reestablished. Normally, the device has now reached the requested track.

On a multitasking, multiuser system like Linux, there are lots of optimizations that can be done here. Look at an output of the command `df` (see output 2.6.1).

So, what benefits does parallelizing bring us? Suppose we enter in `/usr/src`:

```
root@earth:/usr/src/ > tar xzf package.tgz -C /usr/lib
```

Here, `package.tgz` will be untarred into `/usr/lib/package`. To do so, the shell launches `tar` and `gzip` (located in `/bin` and thus on `/dev/sda`) then `package.tgz` in `/usr/src` is read (on `/dev/sdb`). Last, the extracted data is written

```
Filesystem      Size  Used Avail Use% Mounted on
/dev/sda5       1.8G  1.6G  201M  89% /
/dev/sda1        23M   3.9M   17M  18% /boot
/dev/sdb1       2.9G  2.1G  677M  76% /usr
/dev/sdc1       1.9G  958M  941M  51% /usr/lib
shmfs           185M     0   184M   0% /dev/shm
```

Output 2.6.1: Example of a `df -h` Command Output

to `/usr/lib` (on `/dev/sdc`). Using parallelizing, positioning as well as read and write of the disks' internal buffers can be activated at the same time.

This is only one example; there are many more. If this example were a frequent processing requirement, as a rule of thumb, if there are many disks (with the same speed), `/usr` and `/usr/lib` should physically be placed on different disks. Here `/usr/lib` should have approximately seventy percent of the capacity of `/usr`. `/`, due to its access, should be placed on the disk containing `/usr/lib`.

From a certain number of SCSI disks onwards (4–5), consider buying a RAID controller. Thus, operations on the disks are not only quasi-parallel but parallel. Fault tolerance is one of its famous by-products.

Processing speed and size of main memory

The size of main memory is more important in Linux than the processor itself. One reason, if not the main reason, is Linux's ability to dynamically create buffers of hard disk data. Here, Linux uses lots of tricks, such as "read ahead" (getting sectors in advance) and "delayed write" (saving writes until there is a bundle to write). The latter is the reason why you should not switch off your Linux machine. Both items are the reason why Linux is so fast and why the memory seems to fill rapidly. Linux works with shared libraries — lots of programs and applications share the same library. A library call only has to be allocated once. This also means that your memory fills up rather quickly. If you do not know whether to purchase another main board or some more memory, we recommend you do the latter as it increases the speed of Linux.

```
                total  used    free  shared  buffers  cached
Mem:             255    246     9      0      23     44
-/+ buffers/cache:  178     76
Swap:            261     3    257
```

Output 2.6.2: Output of `free -m`

Output 2.6.2 shows that approximately 23 MB of data is being saved in buffers. If you want to access data that is already buffered, this data is almost immediately available.

3 YaST — Yet another Setup Tool

YaST (**yast**) is a very important part of SuSE Linux. YaST installs the system, manages the software (install and uninstall), and provides system administration support. This chapter outlines many important features of YaST, some of which are used during system installation.



Note

This SuSE Linux distribution comes with a version of YaST updated only recently. Therefore, it is possible that the screen shots included in this chapter may differ in detail from what is displayed on your monitor.

3.1 Basic Usage and Keyboard Layout

Start YaST as user `'root'` from the prompt simply by typing the command **yast**.

Navigate the menus and the screens in YaST with the arrow keys and **(Tab)**. Move around in lists with the arrow keys and with **(PgUp)** and **(PgDn)**. Items in the lists can be selected with **(Space)**. Normally, use **(↵)** to confirm your choice. In most cases, close any menus, screens, or dialogs with **(Esc)** key or `'Abort'`. In the more complex dialogs, keep a particular choice by pressing **(F10)** or selecting `'Continue'`. Where you are prompted for a yes or no response, use **(Tab)** to switch back and forth between the two options. If you use YaST in color mode, the active selection is always highlighted in blue.

In rare cases, such as when you are accessing YaST remotely from a non-Linux terminal, the function keys may not be usable. In these cases, the key combination **(Ctrl) + (F) <number>** can be used as a substitute for **(F<number>)**. To simulate **(F10)**, use the key combination **(Ctrl) + (F) (0)**. There are no substitute mappings for the **(F11)** and **(F12)** function keys.

3.2 The YaST Main Menu

When you first start YaST, you will be in the “main menu” (Figure 3.1 on the next page).

'General help for installation' As the title suggests, this menu option provides some general tips on installation.




Figure 3.1: The Main Menu

‘Adjustments of installation ->’ This option opens a submenu for making adjustments to your installation. See Section 3.4 on the facing page for more details.

‘Package management (Update, Installation, Requests)’ This menu option will bring you to the YaST package manager utility which is used to install software packages and remove them from the system (see Section 3.5 page 61).

‘Update system’ Use this for updating the system.

‘system administration ->’ Use this for system administration. See Section 3.7 page 69.

‘Show README file for installation media’ Important late-release information can be found here.

‘Copyright’ Read copyright information.

‘Exit YaST’ Exit the program.

3.3 General Help

This menu gives access to the online help of YaST. Individual topics are organized as hypertext and include a summary of keyboard shortcuts.

3.4 Adjusting Installation

Adjustments to your installation can be made in one of the submenus of ‘Adjustments of installation’ (Figure 3.2). With **(Esc)**, return to the main menu.

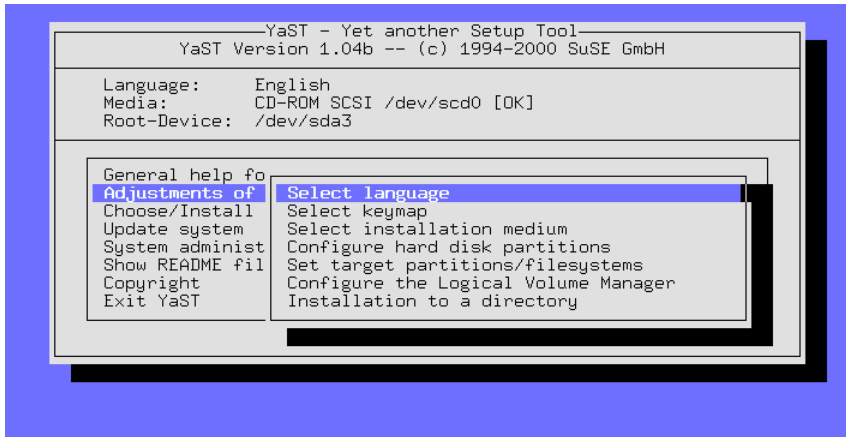


Figure 3.2: The menu for ‘Settings for installation’

The actual installation can only be started once you have specified the target partitions. A number of items can be selected after the installation, useful if you want to update settings.

3.4.1 Selecting the Language

With ‘Select language’, the language of the YaST mask can be changed. The value is stored in the variable `DEFAULT_LANGUAGE` in `/etc/rc.config`. See Section 13.9 page 303.

3.4.2 Selecting Keyboard Layout

With ‘Select keymap’, the keyboard layout can be changed. The value is stored in the variable `<KEYTABLE>` in `/etc/rc.config` (see Section 14.6 page 314). Choose a keytable according to the default language.

3.4.3 Installation Medium

Selecting the menu option ‘Select Installation Medium’ allows selection of the source medium. From this menu (Figure 3.3 on the following page), select the medium from which to install. This selection is made for installing the initial system and for adding packages to an existing system. In another menu, YaST

provides the opportunity to select an installation medium, but this is for applying patches or updates to an existing system. This process is discussed later in this chapter.

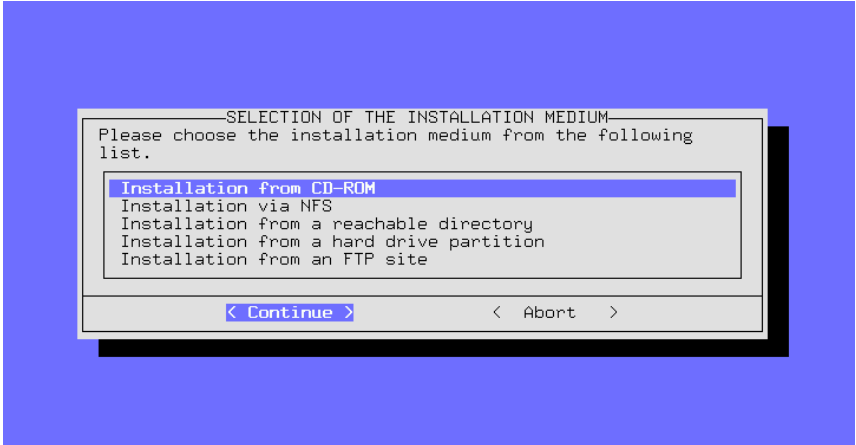


Figure 3.3: Selection of the Installation Medium in YaST

In most cases, you will want to install Linux directly from the CD-ROM. To do this, choose 'Installation from CD-ROM'.

With the menu option 'Installation from hard disk partition', you have the option to install from a hard disk that contains the necessary source files. This option is helpful if your CD-ROM is not supported under Linux (see also Section 3.4.5).

With the menu options 'Installation via NFS', 'Installation from an FTP site', 'Via SMB', install Linux on a system that does not have a CD-ROM drive, but is connected by Ethernet to another computer that does have a CD-ROM drive. For additional information on these options, refer to Section 3.4.6 on the facing page and Section 3.4.8 page 49.

3.4.4 Installation from CD-ROM or DVD

If you plan to install via CD-ROM, specify your CD-ROM drive (see Figure 3.4 on the facing page). If unsure, try 'ATAPI EIDE' drives.

3.4.5 Installation from a Hard Disk Partition

If your CD-ROM drive is not directly supported, you can still install Linux on your system. To accomplish this, however, you will need to use a different installation medium.




Figure 3.4: Selecting a CD-ROM Drive

3.4.6 Installation via NFS

Network installation gives the option of easily performing installations on multiple computers. This is especially useful when only one of the computers has a CD-ROM drive from which the source files can be placed on an NFS reachable hard disk. Before attempting to install via NFS, you should be familiar with how to configure NFS servers.

It is also possible to use this method to install Linux onto notebooks which have a functional PCMCIA network card.

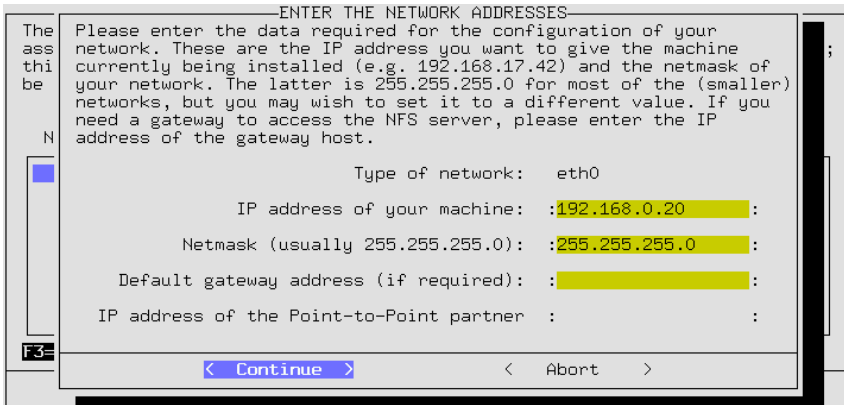


Figure 3.5: Entering Network Data for an NFS Installation

Installation via NFS is available to computers that are connected together via ethernet as well as computers “networked” via parallel ports. This feature is

of particular interest for users who want to install Linux onto laptops and notebooks. If you are planning to attempt such an installation, select PLIP support when booting. Additionally, in this case, you will need to provide additional entries so that the PLIP interface can be correctly configured.

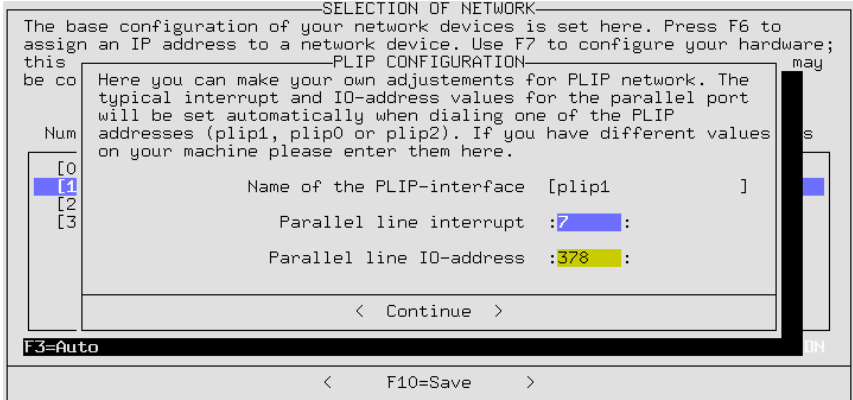


Figure 3.6: PLIP Configuration

If you are directly connected to the NFS server through the parallel port, the IP address of the PLIP partners is the same, as shown in the following examples of the NFS servers. The PLIP interface in most cases will be `plip1`. By “interface” we are referring to the network interface that can be seen and configured with the command `ifconfig`. The hardware parameters will differ only in rare cases.

You should be aware that, by using the parallel port as a PLIP interface, printers will not function on this interface. In many cases, a printer connected to such a port will lead to a permanent reset as soon as the driver starts.

In the following dialog, enter the IP address of the NFS server as well as the directory in which the source files are located. Of course, the NFS server must have exported this directory to make it available to the computer you are installing.

3.4.7 Installation from a Reachable Directory

With this option, it is possible to install additional software later if you already have SuSE Linux running. This is a useful option when you are unable to install using a CD-ROM due to the unavailability of drivers.

To access this drive from YaST, first start YaST. Once this is accomplished, switch to another console. For example, with the `(ALT) + (F2)` key combination you will be presented with the second virtual console. Log in as `'root'`.

>From here, mount the CD-ROM drive *manually* onto an available mount point in your directory tree, for example, with the command

```
earth: # mount -t iso9660 /dev/cdrom /cdrom
```

Now follow through with a normal installation. In the following screen (Figure 3.7 on the next page), enter the directory in which the source files are located.

This should be the directory onto which the CD has been mounted, followed by the directory `suse`.



Figure 3.7: Entering the Source Directory

If, for example, you had mounted the drive or the hard disk partition onto the directory `/cdrom`, use the field as depicted in Figure 3.7.

3.4.8 Installation via FTP

Like the NFS install, it is possible to install SuSE Linux onto a computer that has no CD-ROM at all via FTP. This will work, but only if the basic network configurations have been set up correctly.

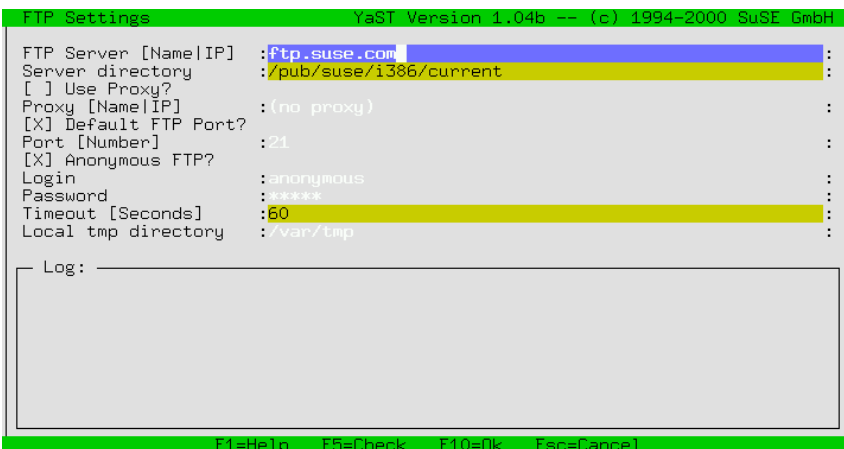


Figure 3.8: Entries for the FTP Installation

- ‘**FTP Server [Name|IP]**’ The name or IP address of the FTP server.
- ‘**Server Directory**’ Where on the FTP server the `suse` directory is located.
- ‘**[] Use Proxy?**’ Only select this if you are sure that you must use an FTP proxy server. Normally, this is *not* necessary.
- ‘**Proxy [Name|IP]**’ This is only necessary if you have opted to use the proxy server.
- ‘**[X] Default FTP Port?**’ This should be selected by default.
- ‘**Port [Number]**’ The default value should be 21.
- ‘**[X] Anonymous FTP?**’ This should always be selected if you want to access a public FTP server.
- ‘**Login**’ In the case where you have *not* selected the anonymous FTP option above, insert a user name.
- ‘**Password**’ Insert a password, if required.
- ‘**Timeout [Seconds]**’ 60 is the suggested value.
- ‘**Local /tmp directory**’ The local directory available for temporary storage.

3.4.9 Partitioning Your Hard Disk

The most critical point of installing a new operating system is partitioning the hard disk. Generally, every operating system uses at least one partition.

Tip

With Linux, it is also possible to install the system on an existing MS-DOS file system, but you should use this option only to “browse”. The performance is substantially less than when installing Linux in its own Linux-specific partitions. The system is not as secure, because there are no file check systems available for MS-DOS and Linux could be influenced while DOS is running.



With Linux systems, you will generally want at least two partitions: one for files and programs and one for memory data swapping which is performed by the running system. You might consider creating multiple partitions for files and programs. The number and sizes of partitions in a Linux system is a question of personal taste and philosophy (see Section 2.5 page 36 and Section 2.6 page 39). A detailed example of partitioning can be found in Section 3.4.11 page 56.

In addition to the partitions you create for Linux, you should also create a swap partition to increase the size of the computer’s virtual memory (see Section 2.6.1 page 39). It is possible to use a swap file instead of a swap partition. For performance reasons, however, this is not recommended. Every time your system

accesses this file, it must do it through the file system. Systems with limited memory should always use a swap partition.

If you have more than one hard disk on your system, you will be prompted to select the drive to partition. Once you have selected the drive, you will see a menu which shows the current status of the partitions on your hard disk (Figure 3.9).

```

-----EDITING THE PARTITION TABLE-----
Fdisk detected the following hard drive geometry:
Disk /dev/sda 255 Heads 63 Sectors 1115 Cylinders.
One cylinder has 8225280 Bytes.

Here you can see the error messages of the fdisk program:

The number of cylinders for this disk is set to 1115.
There is nothing wrong with that, but this is larger than 1024,
and could in certain setups cause problems with:

Current partition table of the selected hard disk:
Device name      From      To      Blocks  Partition type
-----
/dev/sda1        1         1       8001    83  Linux native
/dev/sda2        2        18     136552  82  Linux swap
/dev/sda3        19       592    4610655 83  Linux native
/dev/sda4       593     1115   4200997 83  Linux native

F1=Help  F3=Change type  F4=Delete  F5=Create  F6=View errors

< Continue >          < Abort >

```

Figure 3.9: Partitioning with YaST

The arrow keys \uparrow and \downarrow can be used to navigate through the list of existing partitions. To change the type of file system associated with a partition, for example, from DOS to ext2, use F3 . To delete existing partitions, use F4 . To create new partitions, use F5 .

Changing the type of the partitions (using F3) is necessary, when you create a swap partition. Sometimes, it is not necessary to change the partitioning at all, especially if you decide to use the partition layout from the previously installed operating system. In this case just select the respective partition and press F3 .

Linux has no preferences regarding the type of partition on which it is installed. It is equally acceptable to install Linux on either a *primary* partition or a *logical* partition within an extended partition. The logical partition is a workaround for limitations of modern hardware. It was introduced to overcome the fact that the partition table is only large enough for four entries. If you are interested in having more than four partitions on your hard disk, create one *extended* partition. After the extended partition is created, create multiple *logical* partitions within this one large *extended* partition.



Note

Extended and logical partitions are known to DOS as *extended DOS partition* or *logical drive*, although this concept applies to any operating system.

3.4.10 Setting Target Partitions and File Systems

After you have completed partitioning your hard disk, set how the partitions will fit into the Linux directory tree. Select the menu option ‘Set target partitions/filesystems’.

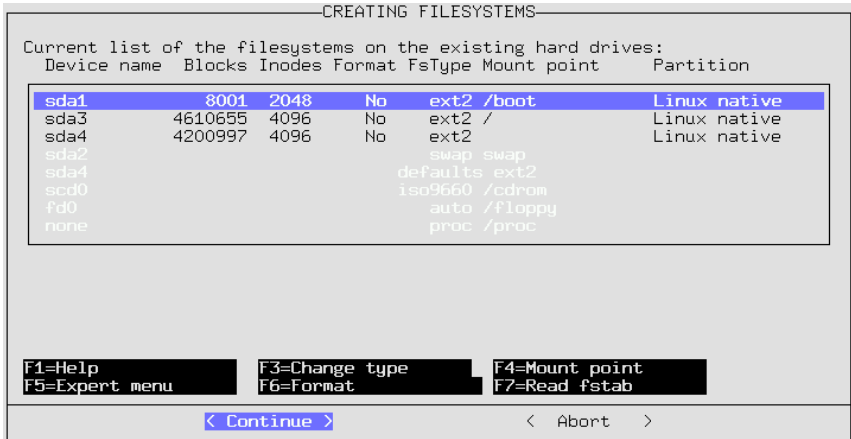


Figure 3.10: Assigning File Systems

In Figure 3.10, you can see an example of partitioning a hard disk. For each individual partition, select if and how it should be formatted as well as at what location in your directory tree it should be mounted.

You must specify one partition as the “root” partition! This partition is necessary because it is the starting point of the entire directory structure. It *is* analogous to the function that roots provide for trees. Make sure you assign ‘/’ as the mount point for your “root” partition.

The menu at the bottom of your screen defines the different function keys that can be used to manipulate the file system. As a rule, you only have this functionality available on Linux partitions. If you highlight partitions formatted for other file systems, the function keys will not be available for use, except (F3), which is used to select the mount point.

Setting the file system type

With (F3), choose between the two file systems, ext2 and reiser.


‘ext2’ — The ext2 file system (*second extended file system*) has been established for many years as the standard file system.

‘reiserfs’ — *ReiserFS* is a file system for the next generation and has recently been included in the official 2.4 kernel. Read `/usr/share/doc/packages/reiserfs/README` if you intend to use it.

In close cooperation with the developer group surrounding HANS REISER and CHRIS MASON, SuSE has extended the high-performance file system to include a “journaling file system”. Through journaling, it is guaranteed that, even for large servers, a file system check can be carried out within a few seconds.

Caution

Do *not* use ReiserFS together with RAID 5 software. For hardware RAID solutions, there are no such limitations. At present, a ReiserFS partition must be at least 34 MB in size.



To ensure there are no problems, create a separate `boot` partition with the `ext2` file system and install `LILO` there. This is the standard SuSE Linux configuration.

The following points should be noted if you want to use ReiserFS for an NFS server:

- `knfsd` generally works, but problems may arise with the extensive renaming of special files (“device files”, “named pipes”) via NFS across different directories.
- `unfsd` functions reliably if the file system in question is exported with the option `no_cross_mounts` or if `unfsd` is started with `--no-cross-mounts`.
- Exporting entire file system hierarchies (possible with `unfsd`) does not function reliably with ReiserFS.


The utilities for ReiserFS are in the package `reiserfs`, series `a`. Note, however, that, as of now, `reiserfsck` is not capable of repairing all file systems. Current information can be found at <http://www.namesys.com/>.

Mount points

The **F4** key can be used to select the location where the selected partition should be mounted on the directory tree. You *must* have *one* partition assigned as the root partition (`/`). If you already have DOS/Windows partitions on your drive, you can also create mount points for them. For example, create a mount point of `/dosc` for the first DOS partition and a mount point of `/dosd` for the second.

Caution

Make sure that you assign absolute path names (starting with `/`) for all mount points and that you create one, and only one, unique mount point for each partition. In addition, never create separate partitions for the `/etc`, `/bin`, `/sbin`, `/lib`, and `/dev` directories. These directories, which contain important commands, libraries, and configuration files necessary for mounting the rest of the file system, must reside on the root partition.



DOS and Windows partitions can be connected to the directory tree as one of three different types.

- As a “normal” **DOS** partition (`msdos`), with all of the normal limitations of a DOS file system
- as a **FAT-Win95** (`vfat`) in which long file names are possible
- as the **UMSDOS** partition, which permits the use of long file names in a “normal” DOS partition. This type is only required, however, if you want to store Linux files on a DOS partition. This is not very efficient, however, if you only need the DOS partition to store data. Whenever possible, you should *avoid UMSDOS!*

Expert Menu for Fine-Tuning the File System Inode Density

Fine-tuning can be performed here in the Expert menu.

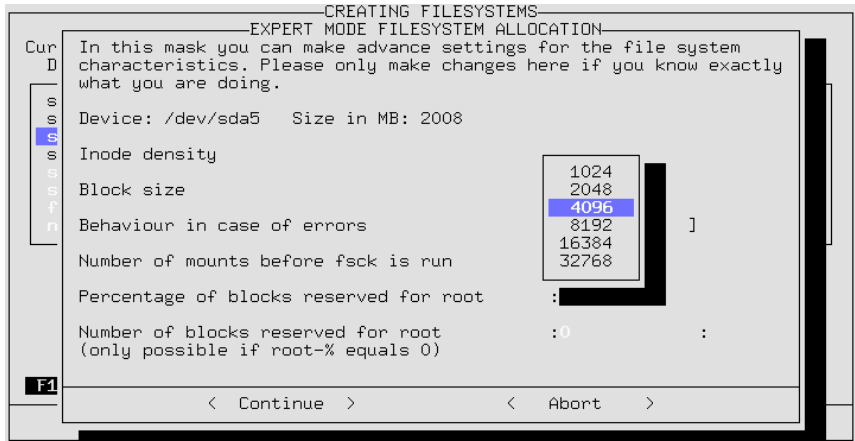


Figure 3.11: Expert Menu for Setting File Systems

The inode density is used to define the anticipated average file size (per inode) on a particular partition. If, for example, you choose a density of 4096 bytes per inode, this means that you expect the average file size to be about 4096 bytes (4 KB).

When a file system is created on a partition, the system computes the number of inodes it will need, using the inode density number and the size of the partition. The number of inodes represents the number of files that can be created on a given partition.


In the example above, if the partition’s size is 4 megabytes, the partition would have less than 1000 inodes. If this partition is filled with many files that are all approximately 1 KB, only one-fourth of the partition could be used before all of the inodes were exhausted. No more files could be added to the partition.

The inode value of 4 KB per file has proven to be a good default value (an excessive number of inodes would result in the unnecessary consumption of disk space, since inodes need disk space, too). If a particular partition is used to perform spooling (necessary for the news service, for example) a density of 2048 bytes per inode is recommended, because news articles are typically very small files.

Deciding on the *block size* has a direct influence on how long a file system check lasts. The suggested block size is 1024 (1 KB) for partitions up to 1 GB, stepping to 4096 (4 KB) for 4 GB partitions. Apart from that, you can specify the conditions for a file system check to be performed. Optionally, it can be taken into account how frequently the file system will be mounted.

Specify the area reserved for `'root'`. For `'root'`, you should always reserve an area if partitions are to be used both by system programs and by regular users. A separate `/home` partition means that space does not have to be set aside for `'root'`.

The *behavior on errors* can be specified. If an error occurs in the file system, operation can proceed as normal (`continue`), the partition can be remounted as read-only (`remount-ro`), or the complete system can be halted immediately (`panic`).



Caution

The option to choose depends on for what your system is used. The default `continue` should remain reserved for experienced system administrators.

Formatting the Partitions

With the `(F6)` key you can select if and how the partitions should be formatted. With brand new disks it is not necessary to check for bad sectors while formatting, but to be on the safe side you can choose to do so. Selecting the formatting and check option will take quite a bit longer than just formatting a partition.

Reading the `fstab` File

The file `fstab` tells Linux what file systems to mount when it loads. If you have Linux already installed on your hard disk, press `(F7)` to read an existing `fstab` file. The mount points listed in the `fstab` file will be automatically included for existing non-swap partitions on your hard disk. Entries in the `fstab` file for other types of file systems (`swap`, `proc`, `nfs-mount`, CD-ROM entries, etc.) will appear grayed out and cannot be changed. They will remain unchanged when you save the `fstab` file.

This feature is used when you are updating your entire base system (see Section 12.1.3 page 272). In such cases, YaST needs to know how your file system and partitions are divided. It is possible to have more than one Linux system installed on your computer.

3.4.11 Configuring Your Hard Disk Manually

Configuring Partitions

What's it all about?

In Section 2.1.3 page 12, you have decided to interactively partition your hard disk. Here we describe the menus. Background information on the various partition types can be found in a document by Andries Brouwer (<http://www.win.tue.nl/~aeb/partitions/>).


Step by step...

This is how to proceed when setting up your partitions:

1. YaST presents a screen divided into several parts (see Figure 3.9 page 51): At the top, you should see your hard disk's parameters. The second part shows warnings and error messages from the program fdisk. Take a closer look by pressing (F6). At the bottom, see the partitions fdisk has found. Here, you should see the partitions of your former operating system and the swap partition, if assigned.
2. To remove existing partitions to use the space for new partitions, do that now. Highlight the partition to remove. Use (↓) and (↑) to scroll through the list. Make sure the correct partition is highlighted. Then press (F4) and confirm by pressing (↵).

Caution

If you remove a partition, all data on this partition will be lost.



If you want to remove multiple partitions, do that now.

3. To use existing partitions of other systems for Linux, do so simply by altering the partition type.

Caution

If you alter the partition type, other operating systems, such as DOS and Windows, may not be able to access data contained on them.



Highlight the partition to change using (↓) and (↑). If you are sure you have marked the correct partition, just press (F3). Now you are presented with a dialog box where you select the partition type. Select either *normal* or *swap* and confirm by pressing (↵).

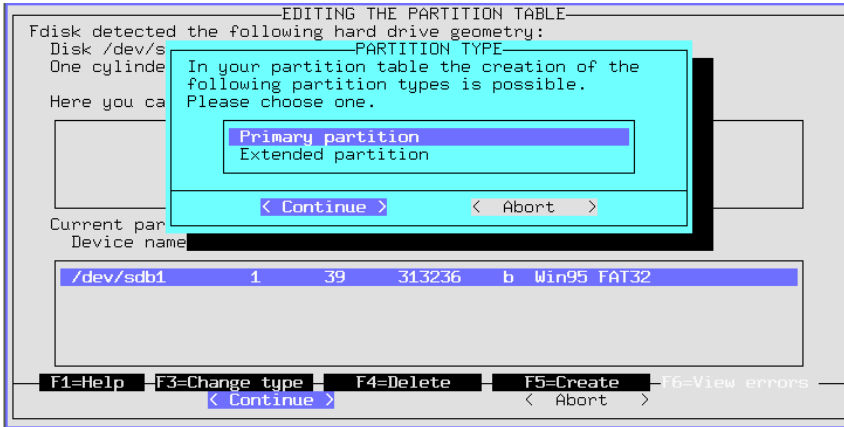


Figure 3.12: YaST — Assigning Partitions

- To assign new partitions, should create them sequentially. Just press (F5). If this does not have any effect, it means that the disk is already full with partitions. In that case, it is necessary to remove partitions (see above) before new ones can be created.


A dialog box opens where you select the type of partition. (Figure 3.12) Select 'primary partition', 'extended partition', or 'logical drive'. Confirm by pressing (↵). Assign up to a maximum of four *primary* partitions. If you need more, you need to assign at least one as an *extended* partition. Within this *extended* partition, you may assign *logical partitions*. See Section 2.5 page 36.

When you have assigned a primary or logical partition, you must enter the device name. YaST displays a list of the available device names. Normally, just select the uppermost entry (e. g., '/dev/hda2') and press (↵).

In the next step, set the size of the partition. YaST recommends starting at the first cylinder of the first free partition. Normally, confirm this default. To set the size, select one of three alternatives: the number of the last cylinder (e. g., 976), the number of cylinders in the partition (e. g., +66), or the size in megabytes (e. g., +100M). Pressing (↵) should now move you to 'Continue'. Here, just press (↵).

Now you should see your freshly assigned partition in the lower part of the screen. If you are not satisfied, remove it by pressing (F4).

- One of your partitions needs to be a swap partition. If there is not already a swap partition, use (↓) and (↑) to select a proper one. Information on the size of a swap partition may be found in Section 2.5 page 36. Press (F3), select 'Linux Swap Partition', and confirm by pressing (↵).
- Is everything to your satisfaction? The drive should be completely populated with partitions. Did you remember the swap partition? Now move the cursor

to 'Continue'.  then continues to the 'Assigning mount points' menu.

Under certain circumstances, YaST checks your swap partition to make sure no errors occurred.

3.4.12 File systems and Mount Points

What's it all about?

In the previous section, all partitions were written to the partition table. In this step, provide additional information for the Linux partitions which have just been configured and also for any DOS or HPFS partitions, if they exist. You have created all necessary partitions and these partitions have been added to the partition table, but they are still not written to disk. There are some vital parameters still missing which will be added in this step.

Summary

This information is partially written to the file system table, `/etc/fstab`, which contains all the relevant information on the file systems to mount. Each entry includes the device, its position in the directory tree, and the type of file system.¹ Furthermore, YaST needs this data to create Linux file systems on the formatted drives. The swap partition is not touched in this step as it has been already created in Section 3.4.11 page 56.

In contrast to the partition table entries (see Section 3.4.11 page 56), the `/etc/fstab` entries are Linux-specific and have no influence on other operating systems.

Also worthy of explanation:

- In Linux, all file systems are linked to one "directory tree". Assign each file system a branch on this tree. This is called its mount point. Even DOS and HPFS file systems can be linked into the Linux directory tree.
- Hard disk space for file systems is administered in Linux by using Inode. An inode is just a small file that points to the corresponding data for a given file. The number of inodes is assigned while creating a file system. If you want to create lots of small files, you will need many inodes, which require some space of their own as well. For file systems containing many large files, fewer inodes are necessary. For more detailed information on this, see Section 3.4.10 page 54.

Step by step

You are now in 'Set target partitions / file systems' (see Figure 3.10 page 52). To determine the file systems of your new partitions, proceed as follows:

¹As well as information for `dump` and `fsck`; see the man page for `fstab` (`man fstab`).

1. First some information about what is to be done:

- For each of your DOS and HPFS file systems (in DOS and HPFS partitions), you *may* create a mount point.
- For each of your new Linux partitions:
 - you *must* select a mount point
 - you *may* select a file system ‘F3=choose type’ (ext2 or reiserfs).
 - you *may* do some fine tuning, with ‘F5=Expert menu’
 - you *may* change the suggested formatting mode ‘F6=Formatting’ (with or without a check).
- The function ‘reading `fstab`’ is normally not needed.

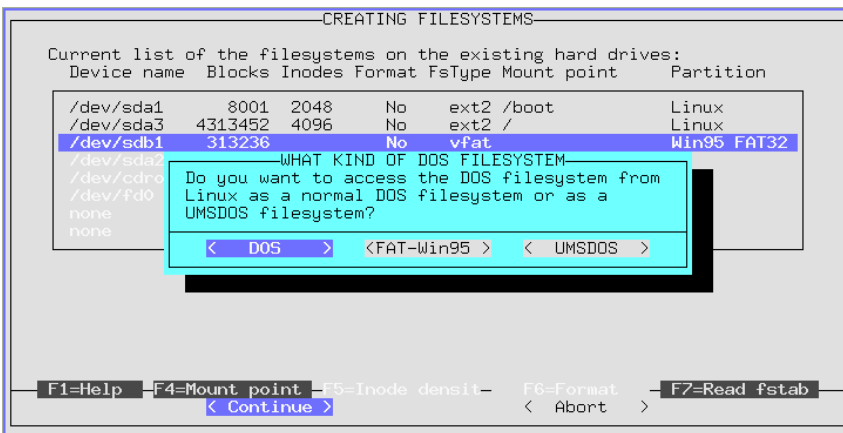


Figure 3.13: YaST — Mount DOS/Windows Partitions

2. Set the type of the Linux file system with (F3). Currently, the tried and tested ext2 and the new reiserfs, with advanced options, are available.

Caution

ReiserFS cannot be used in conjunction with RAID 5 software. For hardware RAID solutions, there are no limitations. reiserfs is very new in comparison to ext2. Bear this in mind if you decide on reiserfs. There is more information about ReiserFS on page 52.

3. If you want to access a DOS or HPFS partition in Linux, select it and press (F4). A window pops up where you can enter a directory for mounting your DOS or HPFS file system. For example, enter /dos (don't forget the slash). Confirm this action by pressing (↵).

If you have selected a DOS partition, another window appears (Figure 3.13). Here, you can choose which method to use to access DOS in Linux (see

Section 3.4.10 page 53). Generally, you will want to use your DOS file system independently and only occasionally access the DOS partition from Linux, for example, to transfer files. For this purpose, choose ‘DOS’ or ‘FAT-win95’ (vfat). UNIX file attributes and long file names, as provided by the UMSDOS file system, are *only* needed if you want to install SuSE Linux on a DOS partition — do this at your own risk!

4. Next, enter the mount points for your Linux partitions. Select your Linux partitions one after the other and press (F4). A window will pop up letting you decide where to mount this partition on your file system tree. Enter the complete path beginning with the slash (/).

Which mount points you choose for your partitions depends on how you plan to organize them. In any case, you *definitely* need a root directory /, which is the “root” for the whole file system. All other file systems and their mount points more or less depend on your preferences. Create a mount point for each Linux file system. You will not be able to access unmounted file systems.

Caution

Directories needed at boot time must be located directly under /. At this stage, the separate branches are not yet put together (mounted). For this reason, you may not assign /bin, /dev, /lib, /etc, and /sbin to other partitions.



5. Ignore the ‘Expert menu’ if you are just starting out with Linux. As a rule, defaults are set automatically to guarantee a secure operation. Only change the suggested values if you really know what you are doing (see page 54).
6. You must now decide whether and how to format your Linux partitions. Since you just created those partitions in Section 3.4.11 page 56, they must now be formatted. If you own a fairly new hard disk, do ‘normal formatting’. If your storage media is not that new, it is better to choose ‘Formatting and checking’.
Select the Linux partitions one after another and press (F6) to select the appropriate formatting. At the end, the menu should resemble Figure 3.14 on the facing page (it is possible you have *not* specified an NFS directory).
7. After selecting ‘Continue’, a dialog box pops up asking whether you are sure. After confirmation, YaST formats your partitions.

3.4.13 Configuring the Logical Volume Manager

Information on configuring the “Logical Volume Manager” (LVM) can be found at <http://www.suse.de/en/linux/whitepapers/> and in the official LVM Howto at <http://www.sistina.com/lvm/Pages/howto.html> The LVM can also be set up with YaST2; to learn more see page 89.

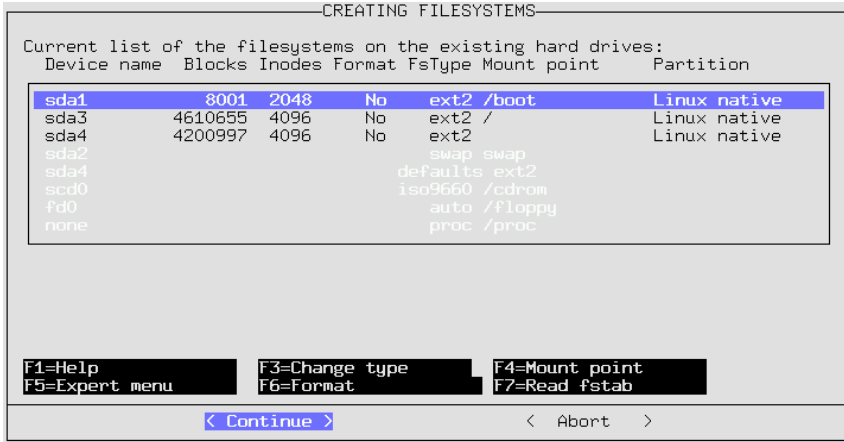


Figure 3.14: YaST — Mount Points

3.4.14 Installation to a Directory

With YaST, you can also make a complete installation into a given directory. Thus it is possible, for example, to update another computer via NFS, to set up a “chroot environment”, or to perform a Linux installation on a second hard disk from the running system. This installation mode is only intended for *special cases*. If certain file systems need to be mounted in the directory specified, you must take care of this yourself before installing with this method.

Apart from that, ensure that the installed system is bootable. Modifications to the `fstab` file created by YaST will also probably be necessary in the newly installed system.

Type, in the entry field, the directory to which to install.

3.5 Choosing and Installing Packages

When you are done configuring your file systems, select ‘Choose/Install packages’ from the main menu to select the packages to install and to start their installation. See Figure 3.16 page 63. At this point, you are also offered the option to define, save, or load your own installation profile (package selection), and to search for packages.

In addition, from this menu, you are able to start installing or to have a look at the implications of your package selection on the space available on your hard disk.

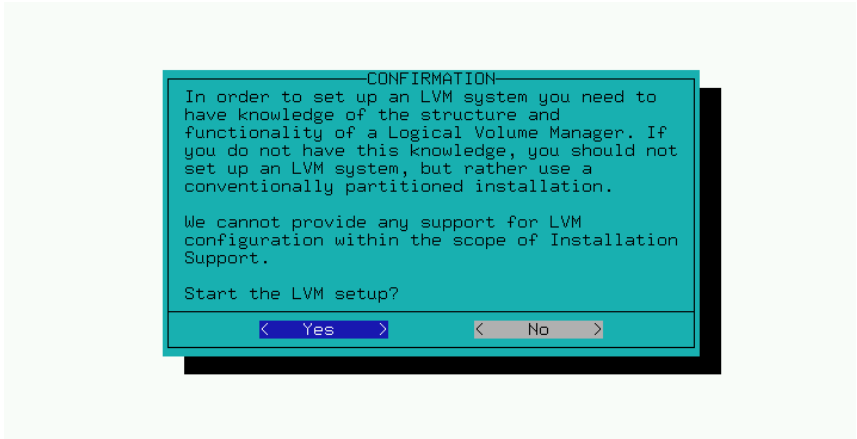


Figure 3.15: YaST — Logical Volume Manager

3.5.1 Load Configuration

There are several predefined configurations listed under ‘Load Configuration’. See Figure 2.10 page 17. Choose any of them using the arrow keys, \uparrow and \downarrow , to move to a predefined “configuration”. With Space , select or deselect an item.

Among them is one labeled “Minimum system”. Choose this if you are running YaST from a floppy disk and have very little RAM. Choosing the minimal system will not prevent you from adding packages after it has been installed. YaST will perform much better after it is installed on the hard disk.

If you have created and saved your own configuration (package selection), use this menu option to load your configuration. This feature is useful if you are installing an identical configuration on many systems.

If you have already installed a system, you should be careful when you load a different configuration, because all packages installed on your system which do not belong to the chosen configuration will be marked for deletion as soon as you select ‘Replace’. If you want to keep these packages, either select ‘Add’ at this point or answer with no later (i. e., after having selected ‘Start installation’) when asked whether you want to delete these packages. If you want to delete most of those packages then you must manually deselect the ones that you do not want to delete. To deselect them, change the ‘[D]’ back to a ‘[i]’ by pressing Space .

3.5.2 Save Configuration

With this menu option, save your current configuration. If YaST is running from a floppy, this will be used to save your configuration. If you have booted from a CD-ROM, you will be prompted to insert a formatted floppy. YaST will save your configuration to this.

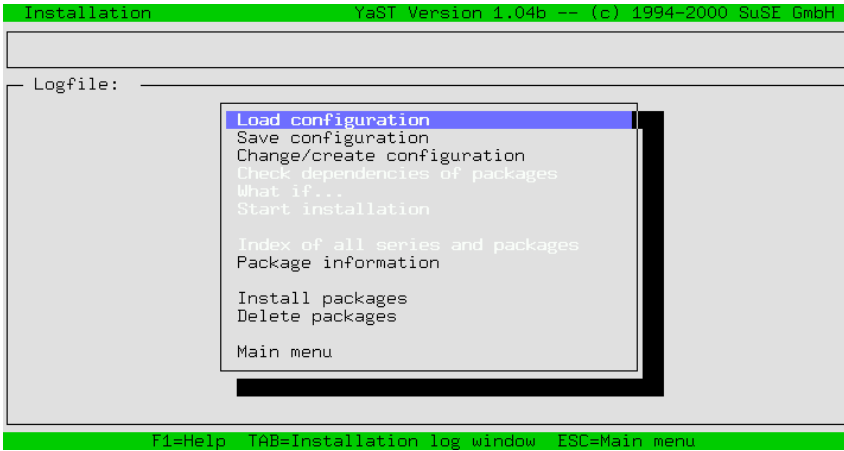


Figure 3.16: YaST Choose/Install Packages — Specifying Installation Size

3.5.3 Changing Your Configuration

If you select ‘Change/create configuration’ you will be presented with a list of series. See Figure 3.17. Edit this list to determine the scope of your installation and change any previously loaded configurations (see Section 3.5.1 on the facing page). After installation, this window can also be used to remove any installed packages.




Figure 3.17: Selecting Package Series in YaST

You can move up and down the list of series with the keys \uparrow , \downarrow , PgUp and PgDn . In the lower window is a description of the current status of the different partitions on your system. The \leftarrow key opens the currently highlighted series

and shows the contents of that series. The **(Esc)** key can be used to exit the package selection without saving any changes. If you have made changes to the package selection, **(F10)** can be used to save those changes and return to the previous menu.

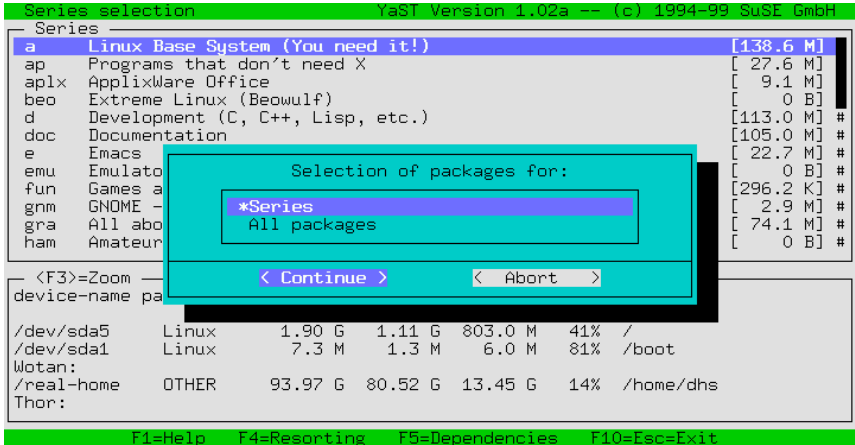


Figure 3.18: Resorting Packages in YaST

The **(F4)** key can be used to select alternate views of the packages (Figure 3.18). You have the choice to view 'series' (which is the default), 'RPM groups', or 'All Packages':

'**series**' The usual categorization according to series.

'**RPM groups**' Order according to RPM group defaults.

'**All packages**' 'All packages (no sources)', 'Base system packages', and 'All sources'.

If you press **(←)** on a series, you will be put into 'package selection' mode for that series (Figure 3.19 on the facing page shows the contents of the series a). If you have loaded a configuration already, the packages belonging to this configuration are marked with an 'x'. In the bottom window, see a short description of the highlighted package. In the right hand window, see the amount of space that is required for your current configuration.

These values are updated as you select and deselect packages. The status windows can be enlarged with the **(F3)** and **(F2)** keys respectively. If you are installing your system for the first time and find that your current partition configuration does not provide enough space for the packages you want to install, go back to the beginning of the installation and recreate the partitions as necessary. If you are trying to add packages to an existing system and find that you do not have enough space, you will need to find another solution (add another hard disk or perhaps delete some unwanted packages).

Each package name is preceded by its current state:

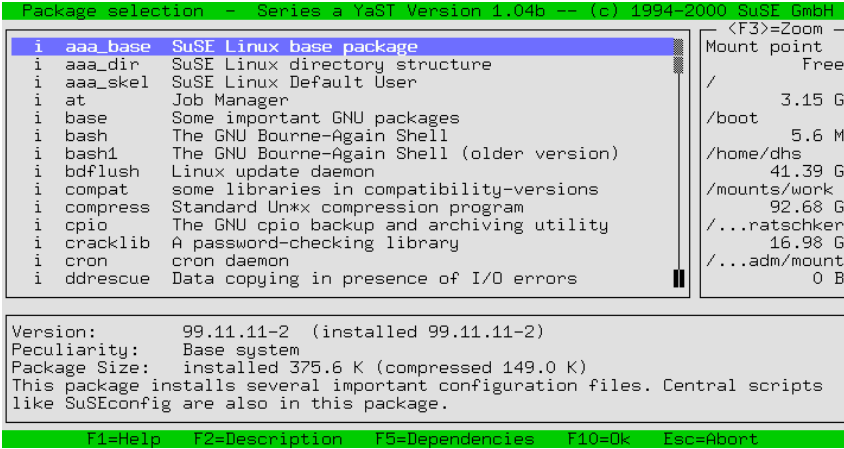








Figure 3.19: Selecting Packages to Install: series a1, Base System

- '[]' indicates that this package is not yet installed
- '[X]' indicates that this package will be installed
- '[i]' indicates that this package is already installed
- '[D]' indicates that this package will be deleted
- '[R]' indicates that this package will be replaced

By pressing , change the status of the package. If the package has not already been installed, pressing  will change the status from '[]' to '[X]' and back again. If the package has already been installed, pressing  will toggle the status between '[i]', '[R]', and '[D]'.  causes the corresponding sources package to be installed as well. For information on source packages, see also Section 12.3.4 page 282. As stated earlier, if you want to save changes made in the package selection, use . If you do not want to save any changes, press .

3.5.4 Checking Package Dependencies

Often the installation of a certain package requires the presence of another package. YaST will check the dependencies of the packages already installed and of any packages selected for installation, subtracting those packages marked for deletion.

Note

If YaST detects that a CD in the drive older than the CD from which the base system was installed or updated, all the dependencies stored on that CD will be ignored.



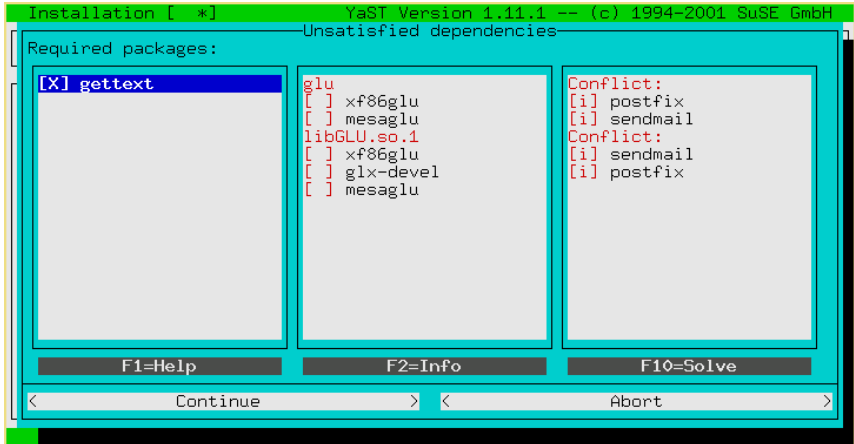


Figure 3.20: Checking Package Dependencies

After performing the check, YaST displays all the dependencies found in a dialog window. See Figure 3.20. There are three different types of dependencies, with each being displayed in its own list. However, if there is no instance of a given type of dependency, the corresponding list will not be displayed.

Required packages. These are packages which are needed by other packages that are installed or selected for installation. By default, required packages are selected for installation. This applies to `gettext` in our example.

Required features. In some cases, installed packages may require a certain feature (rather than a specific package). For example, a package may require the `windowmanager` feature, but no package has been installed so far with this feature. If this happens, the packages matching the `windowmanager` feature will be listed for selection, as is the case with `glu` in our example. At least one of these packages should be selected for installation. Sometimes it may also be useful to select several packages at once from the list.

Conflicts. Sometimes the installation of several packages with similar functionality would lead to unworkable results. For instance, installing then running several mail transport agents at the same time will not give a well-functioning mail system, but, rather, cause a lot of trouble. Therefore, decide in favor of one of the packages in the list and deselect the other ones, as is the case with `postfix` and `sendmail` in our example.

In all these lists, dependencies which are not fulfilled are displayed in red color. This will change to black as soon as your selection has changed to resolve the dependency. To move from one list to another, press `(Tab)`. To make any changes in the list, press `(Space)`. Finally, press `(F10)` to recompute all dependencies and refresh the lists. Your selections are only committed if you leave the dialog by selecting 'Continue'. If you hit `(Esc)` or select 'Abort', the dialog will close, but discard all changes.

3.5.5 What If...

YaST can also be used to review your package selection before committing it. Use this menu item to check that no files would be deleted which are absolutely required for the operation of the system.

3.5.6 Start Installation

This menu option will start the installation of packages. All of the packages selected for installation are copied from the installation medium, decompressed, and written to the appropriate location on your hard disk.

During installation, status messages will be shown in the lower window. The upper window provides information about the current package being installed. After installation, if you want, jump to the installation log window by pressing **(Tab)**. This will allow you to view any messages that may have appeared during the installation of a particular package.

3.5.7 Index of All Series and Packages

This option provides a list of all packages included on the distribution. Those packages marked with an '*' are either already installed or have been selected for installation. This option gives a quick overview of your installation selection.

3.5.8 Searching for Packages

If you cannot find a particular package, but know it is on the SuSE Linux distribution, use the menu option 'Package Information' to find it. You may search through the installation medium, through the entire system, or both. There is even an option for doing a case sensitive search. YaST looks for the string of text entered in the search field and presents a list of relevant packages.

3.5.9 Installing Packages

The 'Install packages' menu option is useful for installing packages not on the distribution. Most importantly, this menu option can be used to update packages directly from our FTP server <ftp.suse.com>, if you have some form of Internet access. The packages on the FTP server are typically bug fixes or newer versions. The 'Install packages' menu option can also be used to install packages directly from your hard disk, a disk, or another source medium. Such packages could include, but are not limited to, packages received from an alternate site, packages you created yourself, or packages downloaded from our FTP server and saved on your hard disk. A few formats are supported: tar archives (.tgz), RPM packages (.rpm, .spm and .src.rpm), and special patch packages (.pat) available, if required, from our FTP server.

The installation procedure is divided into three steps. Press **(F1)** for complete instructions.

- Selection of the installation medium
- Selection of available packages
- Installation of the packages

First, select the package medium by entering 'Source:' and pressing (←). The following options are available: 'Directory', 'FTP', 'Source medium', and 'Floppy'. You may need to change the default path (if you have selected 'Directory' or 'FTP'). YaST will connect to the source medium when you press (←). Thereafter, if you have selected the correct directory, you will be presented with a list of packages to be installed. With 'FTP' it is possible to install directly from the Internet. The address is <ftp.suse.com:/pub/suse/i386/update/7.3> (refer to Figure 3.21). It is not yet possible to install packages via a proxy. You need to have *direct* access to the appropriate FTP server.




Figure 3.21: Installing Packages in YaST via FTP

Tip

If you get an error message, such as "530 User ftp access denied", this means that too many people are currently logged in. Try again later.



First, find the directory that has the `rpm` package to install. Then select the packages to install with (Space). Once the package has been selected for installation, press (F10) to install it. The package is first copied to the `/tmp/ftp<processID>` directory then installed. If something goes wrong, you can still install it manually (see Section 12.3.2 page 279).

3.5.10 Deleting Packages

Selecting ‘Deleting Packages’ provides a list of the packages currently installed. This list includes “foreign” packages — packages which are not part of the SuSE distribution. YaST cannot update these packages or check for any dependencies, unless they were installed in an RPM format.

It is easiest to replace these packages with packages from the SuSE distribution. Just select the packages to replace. A short description is displayed for each package if you press (F2), but only if that package contains description information. Press (Space) to select a package to delete and (F10) to delete it. Afterwards you can reinstall those packages from the SuSE CDs.

3.6 Updating the System

Only use the menu item ‘Update System’ if your base system is a fairly recent one compared to the source medium. If there is doubt, YaST will tell you in no uncertain terms. Begin a system update as described in Chapter 12 page 271. Individual packages can be updated with YaST as explained in Section 3.5.9 page 67.

3.7 System Administration

In addition to assisting with the installation, YaST is a powerful tool for performing system administration.

Select ‘System administration’ from the main menu (see Figure 3.22).

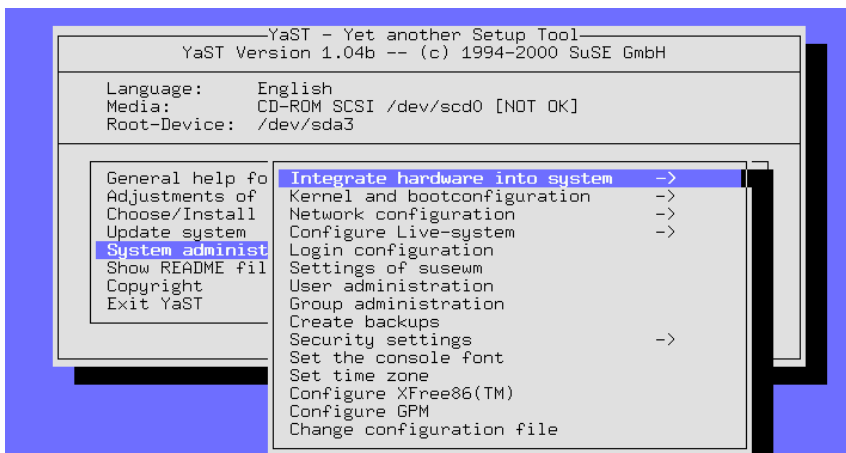


Figure 3.22: Administering Your System

From the main menu in YaST (see Figure 3.1 page 44), choose the menu option ‘System Administration’ to access the system administration features of YaST (Figure 3.22 on the page before).

3.7.1 Integrating Hardware into the System

With this option, specify the hardware used. In most cases, YaST will create a link from the standard device to your particular hardware device. This will make it easier for you to access this device without having to remember the exact name of the device.

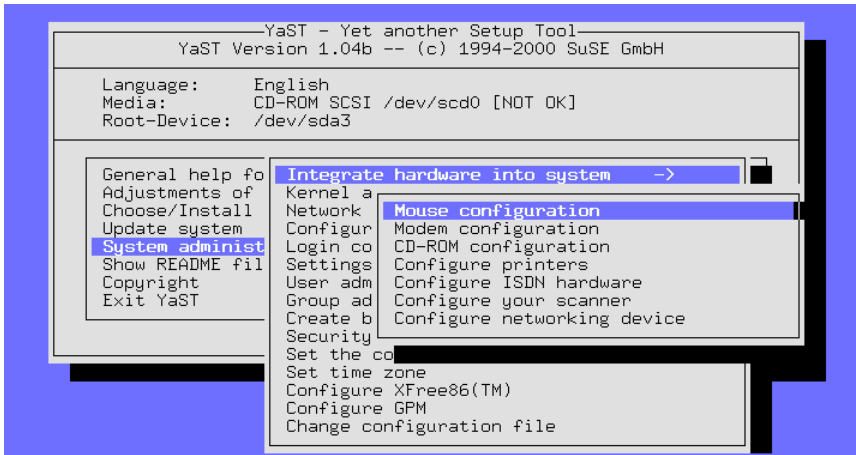


Figure 3.23: Hardware Configuration

Printer configuration is a bit complicated. This is described in the following section.

Printer Configuration

Accessing a printer under Linux is not trivial. The technical background is described in detail in Chapter 9 page 197.

YaST provides an interface for easily setting up `apsfilter` for your printer (Figure 3.24 on the next page).

Under the menu option ‘Printer name’, you will be presented with a list of printers that are supported by Ghostscript. Select your printer from this list. If your printer does not appear on this list, select ‘Other printer’ then the option ‘<userdefined>’. It opens a screen where you can enter your printer type. The command `gs -h` will display a list of printers supported by the currently installed version of Ghostscript. If your printer appears on this list, enter it as the ‘<user-defined>’ printer type. The configuration of `uniprint` drivers can

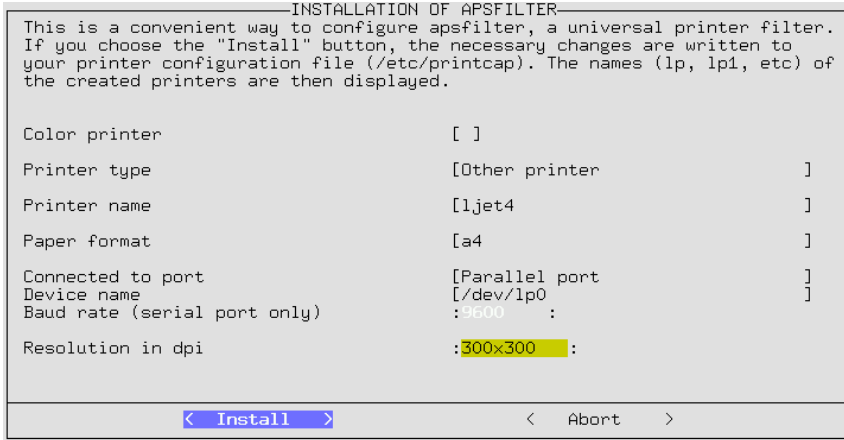


Figure 3.24: Setting Up apsfiler with YaST

be accomplished with YaST2 or the `SETUP (lprsetup)` program. This program is part of the `apsfiler` package. See Section 9.5 page 208 for more information.

Most normal printers are connected to the parallel port (printer port) or USB. In most cases `/dev/lp0` will be the correct port — the first parallel port.

Tip



If in doubt, first set the following values in the BIOS for the parallel port:

- IO address 378 (hexadecimal)
- interrupt 7
- as mode: normal (or SPP)
- and switch off DMA (should be switched off in normal mode)

These are the normal standard values.

The settings for the menu option ‘Resolution in dpi’ should be verified. If you are unsure, refer to your printer documentation.

3.7.2 Kernel and Boot Configuration

This set of submenus can be used for your boot and kernel configuration:

Selecting a boot kernel

Install one of the available precompiled kernels (from the SuSE CD) with the menu option ‘Select boot kernel’, if you want a different kernel than the

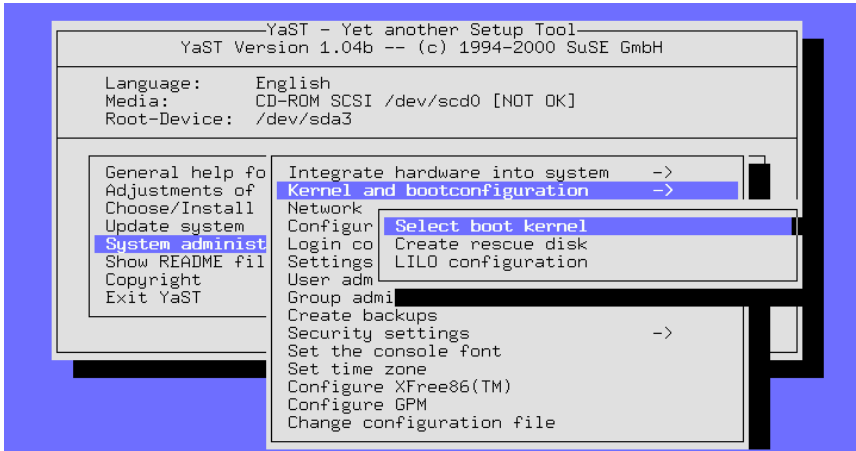


Figure 3.25: Kernel and Boot Configuration

one chosen during installation. In addition to installing the kernel selected, YaST will also copy the kernel configuration file (`.config`) to the kernel source directory (`/usr/src/linux`).

Select the appropriate kernel for your computer. Then, as a rule, you should have LILO installed again, using YaST, if you have already booted successfully with LILO. Otherwise, see below in the guide on configuring LILO.

Advanced Linux users might want to consider recompiling their own kernel (see Chapter 10 page 229). One advantage of recompiling the kernel is that it can be customized to fit your exact hardware configuration and can be made smaller and faster. If you are new to Linux, it is *highly* recommended to use the precompiled standard kernels. You are qualified for “installation support” only if you are using the standard kernel. In other words, by recompiling your own kernel, you no longer qualify for “Installation support”.

Creating a rescue disk

An “emergency” or “rescue disk” is also a good idea, in case your computer will not start at all. See Section 13.6 page 297 for details on how to use this floppy disk.

Configuring ‘LILO’

YaST also provides a front-end for LILO (Linux Loader) to ensure correct configuration (Figure 3.26 on the next page). LILO can be used to boot other operating systems, such as OS/2, DOS, or Windows 95/98, but with Windows NT you should be careful. Additional information on the individual configuration options is available in Chapter 5 page 99.

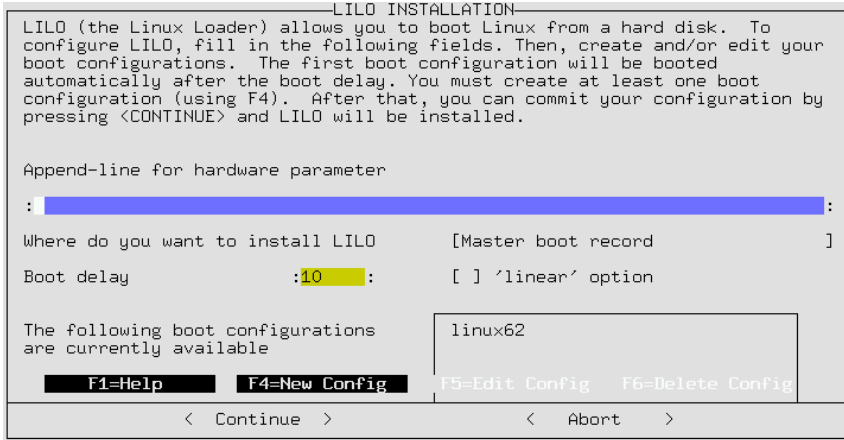


Figure 3.26: LILO: Installation

The ‘append line for kernel parameters’ usually remains empty unless you already need to use additional parameters for booting. Only if this is the case should you enter these parameters here, but *without* the preceding kernel name `linux`. Details can be found in Section 5.4.2 page 107.

‘Where do you want to install LILO?’: If Linux is the *only* operating system on your computer, Master boot record is the correct choice. If you are using “another” boot manager, select Boot sector of the `/boot` partition. If you have not created a `/boot` partition, select boot sector of the root partition. The choice On floppy disk installs LILO on a floppy. The technical background of these choices is given in Section 5.3 page 103.

‘Boot delay’: This entry is in seconds. 10 seconds is a reasonable value to use. “linear” Option’: In most cases, this option is *not* necessary. For additional information, see Section 5.4.2 page 107.

The ‘F4=New Config’ menu option can be used to create a “new configuration”. It is recommended to use the standard configuration `linux` as the configuration name. To edit an existing configuration, use the ‘F5=Edit config’ menu option. When using either of these keys, YaST will present a configuration screen, as shown in Figure 3.27 on the following page. The individual options on this screen are described below.

‘Configuration name’: Choose what you wish, but `linux` is a good name for your first configuration.

‘Which operating system’: Your three choices are Boot `linux`, Boot `DOS/Win` (this option is for Windows 95/98 as well), and Boot `OS/2`. ‘(root-) partition to boot’: This detail is predefined. As a rule, you should not change anything here. Press **F3** and choose the appropriate partition.

‘Kernel optional’: Only select this option if the kernel is *not permanently* available. This would be chosen if you wanted to boot to a test kernel only once.

‘Kernel to be booted by Lilo’: Here, as well, you do not need to change

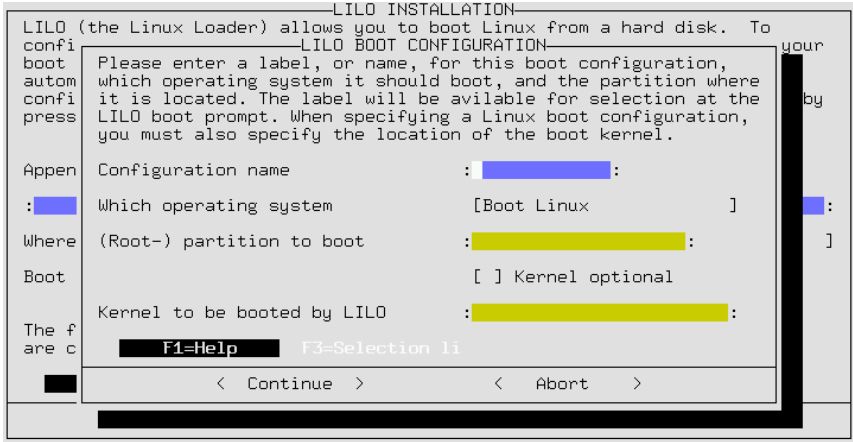


Figure 3.27: LILO: Boot Configuration

anything. The usual location in SuSE Linux is `/boot/vmlinuz`. By pressing **F3**, browse through the directories and select the kernel. With ‘Continue’, LILO is installed. With ‘Cancel’, stop the installation of LILO.

Tip

If you are planning on compiling your own kernel, we recommend that you create a second configuration in LILO. The purpose of this new configuration is a backup option in case your newly compiled kernel does not boot. Create a new configuration and call it `old`. Configure it to use the `/boot/vmlinuz.old` kernel. Copy the current kernel `/boot/vmlinuz` to `/boot/vmlinuz.old`. When configuring this option, you should also enable the ‘Kernel optional’ option. Now, in the event that you have compiled a new kernel (see Chapter 10 page 229), which for some odd reason will not boot, you still have a way to boot to your original kernel.



3.7.3 Network Configuration

General network configurations can be done with YaST as well. See Figure 3.28 on the facing page. Even if you are not connected to a network, it is a good idea to configure your network. Many programs depend on a correctly configured network to function properly.

‘Network base configuration’ This is where you can configure the IP addresses for the appropriate network interfaces (network card, PPP, ISDN, etc.) (Figure 3.29 page 76).

Select the services DHCP or BOOTP via ‘F3=Auto-IP’ to use an automatic IP configuration. Using ‘F6=IP-Address’ gives you the ability to change

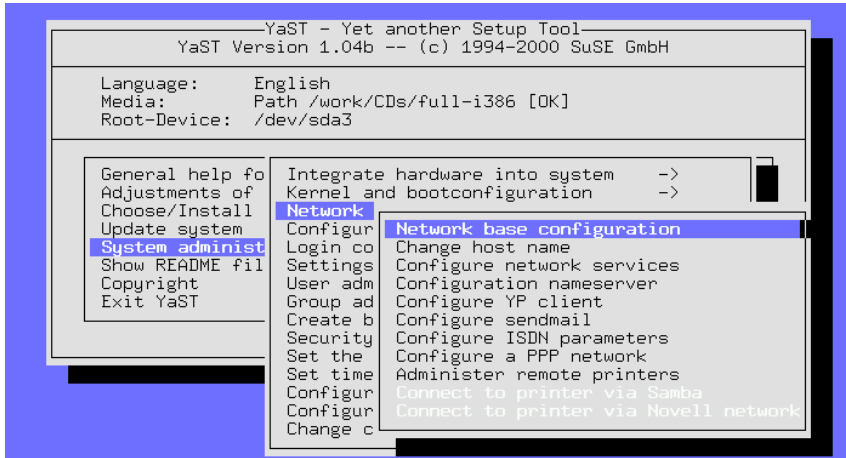


Figure 3.28: Network Configuration

the ‘MTU’. If in doubt, keep the entry blank or use the provided default value. MTU is the “Maximum Transfer Unit” and sets the maximum size of packages used on that network device. Sometimes, your provider forces you to use a special value for ‘MTU’. The standard value, and maximum, is ‘1500’. If the connection to the provider is poor, reduce the value of ‘MTU’.

‘Change host name’ The hostname and domain name of your computer can be set here.

‘Configure network services’ With this option, configure basic network services, such as `inetd` (for telnet, ftp, printing, etc.), `portmap` (for NFS server and NIS), and `rpc.nfsd` (for NFS server).

‘Configuration Nameserver’ It is possible to specify one or more name servers.

‘Configure YP client’ This option will only be active if you have installed the package `ypclient`, series `n`.

‘DHCP Client’ This option will only be active if the package `dhclient`, series `n` has been installed.

‘Configure Sendmail’ With this option, it is possible to install a configuration file for sendmail. There are a few configuration files available that will suit most needs.

‘Administer remote printers’ With this option, configure your machine to access a printer on a TCP/IP network. The configuration options in this screen are relatively self-explanatory. ‘Name of printer’ refers to the name by which the local machine refers to the printer (a good choice would be `remote`). ‘Spool directory’ will be automatically configured according to the name given above. ‘Server name’ is the IP address or the name of

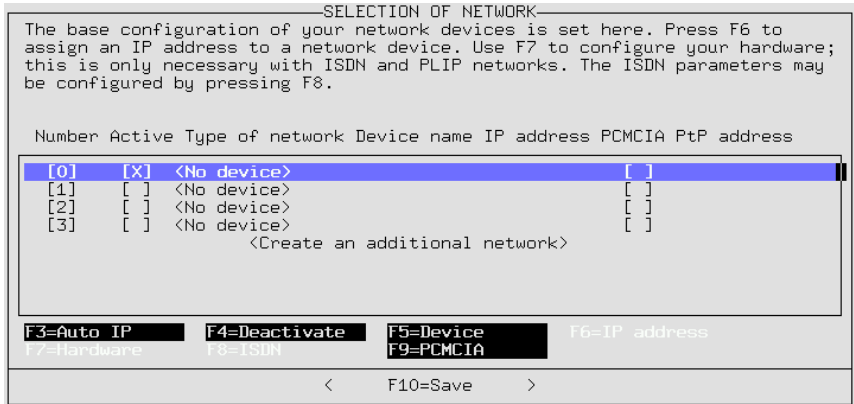


Figure 3.29: Network Base Configuration

the print server. ‘Name of printer’ is the name by which the print server refers to the printer (in most cases this will be `lp`). If you need a prefilter for this printer, refer to Section 9.7 page 215.

‘Connect to printer via Samba’ This option will only be active if the package `samba`, series `n`, has been installed. With this menu item, it is possible to connect to a printer connected to a Windows computer.

‘Connect to a printer via Novell network’ This option will be active only if the package `ncpfs`, series `n`, has been installed.

‘Configure ISDN parameters’ This option will only be active if the package `i4l`, series `n`, has been installed. Follow the detailed menus available.

‘Configure a PPP network’ Here, PPP can be set up.

3.7.4 Login Configuration

Set whether the system will boot to a text mode or to a graphical mode with the X Window System. If you want to use the graphical login, choose among the rather simple XDM, KDE’s full-featured KDM, GNOME’s login manager GDM, and others. If you choose KDM, also select which users are permitted to **shutdown** the system. The other login option is via the text console. With this option, start the X Window System with the command `startx` (Figure 3.30 on the next page).

Note

If you are not absolutely sure that the X Window System is configured correctly, you should not activate either of the graphical logins. Test whether you can start the X Window System first, from the console, before activating the display manager.






Figure 3.30: Login Configuration

If you do configure your system to start with a graphical login, the file `/etc/inittab` will be configured to set runlevel 5 as the default runlevel (see Section 19 page 321). However, if runlevel 5 has been configured on your system to serve a purpose other than starting up a display manager, *none* of the display managers will be started.

3.7.5 Settings for susewm (the Window Manager)

With this menu option, set the “Graphical Desktop”.

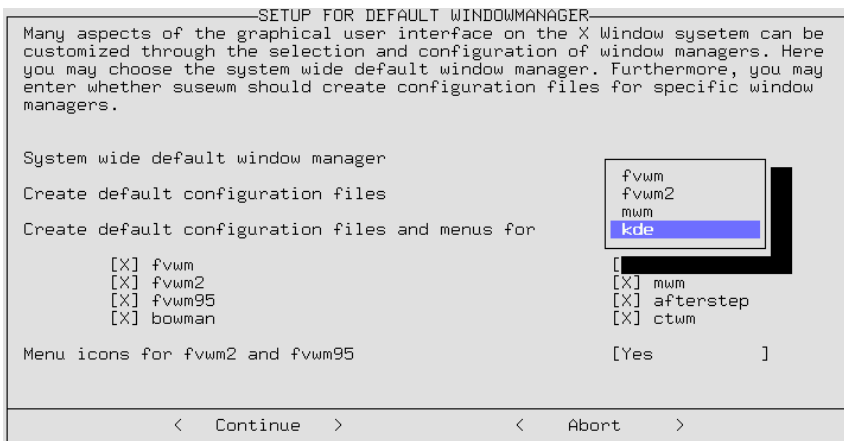


Figure 3.31: Setting Up susewm

Select your default window manager. You could also tell susewm which configuration files it should create and maintain.

3.7.6 User Administration

YaST provides a convenient means of creating, deleting, and maintaining user accounts. The menu option for managing users is titled 'User Administration' and is a submenu of 'System Administration'.

```

USER ADMINISTRATION
In this mask you can get information about existing users, create new users,
and modify and delete existing users.

User name                :tux :
Numerical user ID        :501 :
Group (numeric or by name) :users :
Home directory           :/home/tux :
Login shell              :/bin/bash :
Password                 :***** :
Re-enter password        :***** :
Access to modem permitted [ ]

Detailed description of the user
:tux, the penguin :
F1=Help      F3=Selection list  F4=Create user
F5=Delete user F6=Password times  F10=Leave mask

```

Figure 3.32: User Administration with YaST

Here, you can enable 'Access to Modem'. If you do this, the user is added to the groups 'uucp' and 'dialout'. All members of these groups are able to connect and disconnect PPP connections.

When new user accounts are created, default configuration files are copied from the `/etc/skel` directory into the home directory of the new user. This is a process by which the personal configuration files can automatically be created. Users are free to change those configuration files in their home directory as they wish, of course.

When creating or deleting user accounts, YaST runs the following two scripts:

- When creating a user account, the script `/usr/sbin/useradd.local` is run, if it exists. This script will create the necessary entries in the `/etc/passwd` and `/etc/shadow` files. The home directory is created and the files in `/etc/skel` are copied to it.
- Before deleting a user account, the script `/usr/sbin/userdel.local` is run. The relevant lines in `passwd` and `shadow` are removed. The home directories will remain.

Both scripts can take the user name as a parameter. If additional information is needed (user ID, login shell, home directory), it can be found in the `/etc/passwd` file.

If you are an advanced user and do not want to use YaST to create user accounts, you have the programs `useradd` and `userdel` at your disposal.

Notes on PAM

The configuration files on PAM, *Pluggable Authentication Modules*, can be found in `/etc/pam.d`. Documentation for programmers and system administrators is located in `/usr/share/doc/packages/pam`.

SuSE Linux is able to handle MD5 passwords. With MD5 encryption, passwords can be longer than 8 characters (up to 128 characters). Since MD5 encryption is *not* compatible with the standard Unix `crypt()` function, most commercial Unix systems and some programs do not work with MD5 passwords. Be careful if you enable this feature. Notes on configuration can be found in `/usr/share/doc/packages/pam/md5.config`.

3.7.7 Group Administration

YaST can also be used to create user groups.

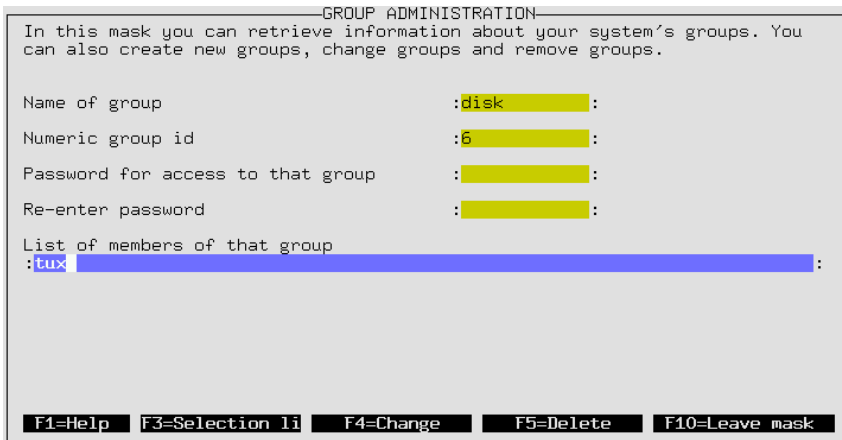


Figure 3.33: Group Administration with YaST

Under Linux, as with UNIX in general, each user must be assigned to at least one group. This is necessary because permissions for certain files depend on to which group the user belongs. User groups can be used to restrict or allow users access to certain directories. Group passwords can also be used to restrict or allow access to certain directories.

Some groups already exist under Linux, for example the user groups `'users'`, `'root'` and many more. The group name, such as with `'users'`, is just a textual representation of the group. Linux recognizes the groups according to the "GID" (Group ID). The configuration file for groups is in `/etc/group`.

This is just background information that most regular users do not need to know. With YaST, it is easy to create groups. In the menu 'System Administration' there is a submenu for 'Group Administration'. This screen is illustrated in Figure 3.33 on the page before.

3.7.8 Creating Backups

This option helps you to back up all modified and new files and packages to a file or tape. These are configuration files in most cases.

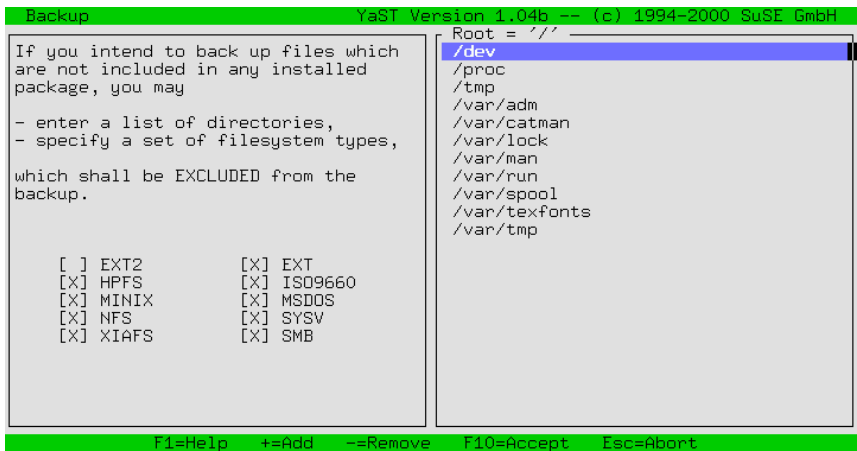


Figure 3.34: Backup with YaST— Choosing Directories to Exclude

The dialog consists of three parts (see Figure 3.34):

1. Choosing the files to back up:

Here, tell YaST which directories should be excluded from the backup. Pre-defined are `/tmp`, `/dev`, and `/proc`. Add mounted CD-ROMs or NFS-mounted file systems to this list. The less you want to be backed up, the faster it will run, since unnecessary comparisons with package lists are omitted. Using `(+)` and `(-)`, add new directories or remove them. Pressing `(F10)` leads to the next step.

2. Searching:

In this step, YaST searches for files which should be backed up. The number and size of the packages found are updated while searching. After this has been done, there will be a list with all the files that have been found. Here you can still deselect files using `(Space)`.

3. Entering commands:

Decide how those files are going to be saved. You can give archive names, options, and more.

This back up mechanism can only work if the dates of the files have not been otherwise changed. Furthermore, this function requires considerable RAM. File names of an ordinary CD take up to 6 MB RAM. Also, you need enough free disk space to save the backup archive. Compressing the archive will lead to a file reduced in size — approximately half of the original. The best way to do backups is to use a tape.

3.7.9 System Security Settings

Only change the defaults if you have detailed knowledge of the measures in question.



Note

If you do not want to allow any 'root' logins via ssh, `/etc/ssh_config` must be adjusted. In `/etc/sshd_config`, the default is set to **PasswordAuthentication yes**. If ssh is correctly configured, 'root' logins are secure across a network.

3.7.10 Setting the Console Font

Here, we will focus on the screen for the font settings in the Linux console. Linux enables you to load various fonts for the console. In this settings screen, try out the different available fonts and apply one as your standard font. If you select 'don't load font', the font saved on your graphics card will still be applied.

All of this, of course, does not have any effect on the fonts in a graphical interface such as KDE or GNOME. See the results in YaST if you have started YaST directly from the ASCII console, but not if you have started it from a kterm or an xterm.

3.7.11 Setting the time zone

The time zone has been preset during the SuSE Linux installation, but you can change or correct these settings later. Simply move the blue bar using the cursor keys \uparrow and \downarrow to the desired selection and then confirm with 'Next'.

Afterwards, a window will appear which will ask you if you want to set your BIOS clock to the local time or leave it set to GTM (Greenwich Mean Time).

3.7.12 Configuring XFree86(TM)

The X Window System (XFree86) can be configured with different configuration tools. First try to use SaX. SaX is described in more detail in Section 6.3 page 132. The technical details can be read in Chapter 6 page 129.

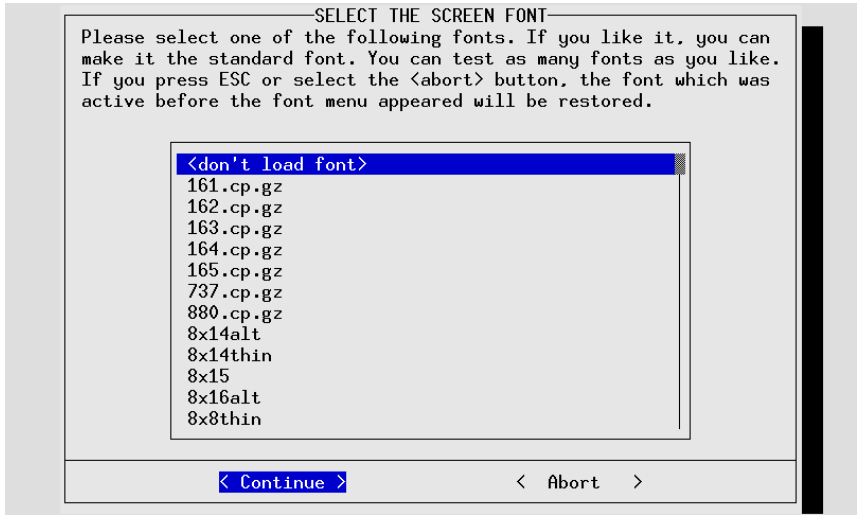


Figure 3.35: Setting the Console Font

3.7.13 Configuring gpm

This screen controls whether to load the mouse support and with which parameters to start it. This allows you to cut and paste text between several virtual ASCII consoles. Again, this entire process is only relevant for ASCII consoles, since the mouse is controlled in the graphical interface by X11 and thus not affected by the gpm settings.

3.7.14 Changing the Configuration File

SuSE Linux is maintained by one central configuration file, `/etc/rc.config`, and the out-sourced files in the `/etc/rc.config.d` directory. These files are read at boot time by the boot scripts that configure your system.

It is possible to change individual variables in these files with YaST. By doing this, you can make configuration changes to your system without having to know every detail about all the special configuration files that have been affected by these changes.

The menu option for changing the configuration file is under the 'System Administration' in the menu titled 'Change Configuration file' With the cursor, select the variable to change and press `(←)` or `(F3)`. After you have used YaST to change one of the values, the script `SuSEconfig` will run automatically.

If you decide to manually change this value by editing the file `/etc/rc.config`, remember that you must manually run the script `SuSEconfig`. This script will do the necessary updates to the relevant configuration files, according to the values that have been set or changed in `/etc/rc.config`. Additional details on this subject can be found in Section 14.6 page 313.

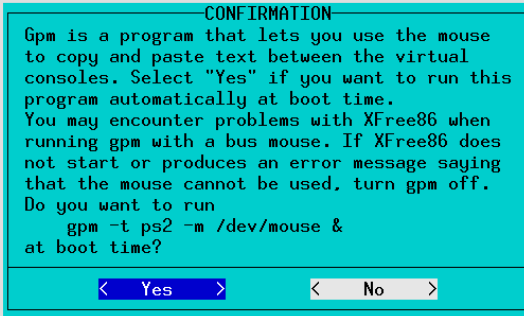


Figure 3.36: Configuring gpm

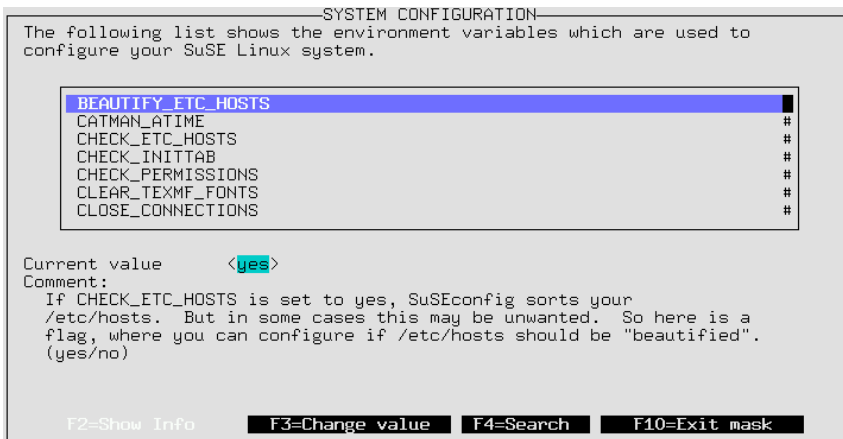


Figure 3.37: Editing the Configuration File with YaST

4 YaST2 — Configuration Tools

YaST2 assists in extending your SuSE Linux system with additional hardware components, such as a printer or sound card, configuring and installing system services, Internet access, and software, or deleting undesired packages. The “Configuration” manual contains a description of various ways of accessing YaST2.

The YaST2 Control Center

After starting YaST2, the YaST2 Control Center will open. The left window panel is divided into the categories of Hardware, Network/Basic, Advanced Network, Security & Users, Software, System, and Miscellaneous. When clicking on an icon, the respective contents will be listed to the right. Configuration usually takes place in several steps. YaST2 guides you through all the dialogs with ‘Next’. A help text for each topic is shown in the left margin of this window, explaining the necessary entries. Once the required data is entered, complete the process in the last configuration dialog with ‘Finish’. Then, the configuration is saved.

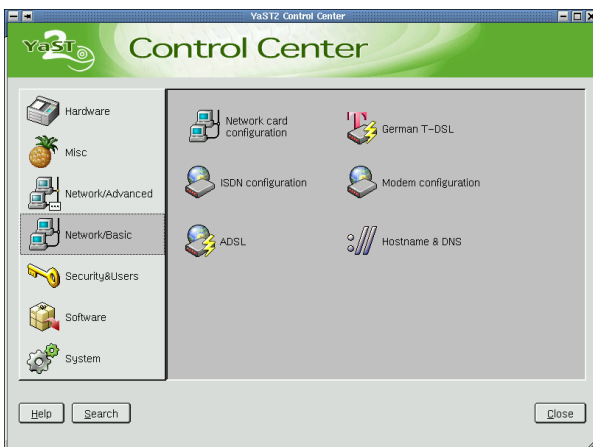


Figure 4.1: YaST2 System Configuration and Administration

4.1 Hardware

New hardware must first be installed and connected according to the vendor's instructions. Connect the external devices such as printer or modem and start the corresponding YaST2 module. The majority of conventional devices will automatically be recognized by YaST2, at which point, the technical information is displayed. If autodetection fails, YaST2 will present a device list (model, manufacturer, etc.) from which to select the appropriate device.

Under 'Hardware', find the configuration tools needed for configuring one or more printers and sound cards. Refer to the hardware information for data pertaining to the hardware autodetected by YaST2. This is especially useful for support requests. The module for configuring the graphical interface (X11) should already be familiar from the installation. It can also be used in the YaST2 Control Center to connect another monitor to an already running system. The same is true for the mouse and keyboard. Finally, there are more modules for configuring scanners and TV cards under 'Hardware'. Descriptions of the different YaST2 tools can be found in the "Configuration" manual.

4.2 Network/Basic

Under 'Network/Basic', YaST2 provides basic configuration tools to pave your way into the Internet: configuration of ADSL, T-DSL (in Germany), network card, ISDN, modem, host name, and DNS. Documentation for this is also located in the "Configuration" manual.

4.3 Network/Advanced

For advanced Internet users and network administrators, there are modules for starting and stopping system services such as inetd, Sendmail (with expert configuration), NFS client and server, routing, expert networking, and NIS client. Instructions are found in the "Network" manual. 'Network for Experts' offers the same functionality as 'Network card configuration' under 'Network/Basic', but you can still configure your modem and ISDN here.

4.4 Security and Users

Here, YaST2 offers tools for simple user and group administration. The configuration module 'System Security' provides various preconfigured levels and expert setting options. Read more about this in the "Configuration" manual.

4.5 Software

To install or delete software, use this module. Use it to change the installation medium. Also, there are two update tools: one for the "normal" system update

and one for the online update which runs over our FTP server. Information on this can be obtained in the “Configuration” manual.

4.5.1 Patch CD Update

As opposed to the online update, the patches are not retrieved from the FTP server, but instead scanned from CD. This CD is available for “SuSE Linux Enterprise Server” customers. The advantage is that it goes a lot quicker with a CD.

After the patch CD is inserted, all the patches stored on the CD will be read into and displayed in this YaST2 module screen. Select which one to install from the patch list. If you forgot to put the CD into the drive, a warning will appear. Then insert the CD and resume updating the patch CD.

4.6 System


Under ‘System’, you will have the option of reconfiguring the boot mode, creating a boot or module disk, or setting the language and time zone. In the “Configuration” manual, find all the necessary information.

4.6.1 RC.Config Editor

`/etc/rc.config` is the file where the most important SuSE Linux system settings are stored. The RC.Config editor presents the settings options in an easy-to-read manner. The values can be changed and subsequently applied to the file. In general, it is not necessary to manually edit them, however, because the file will automatically be adjusted when installing a package or configuring a service.

Caution

Do not modify the `/etc/rc.config` file if you lack previous experience and knowledge. It could do considerable damage your system.



4.6.2 Runlevel Editor

Runlevels in Linux are described in more detail in Section [14.2 page 307](#).

After this expert module starts, it will be initialized. The current default runlevel will be shown in the next dialog. This “operation mode” will start after your system boots. In SuSE Linux this is usually runlevel 5 (full multiuser operation with network and KDM, the graphical login). Runlevel 3 also works well (full multiuser operation with network). At this point, with the help of YaST2 a different default runlevel can be set. See Table [14.1 page 308](#).

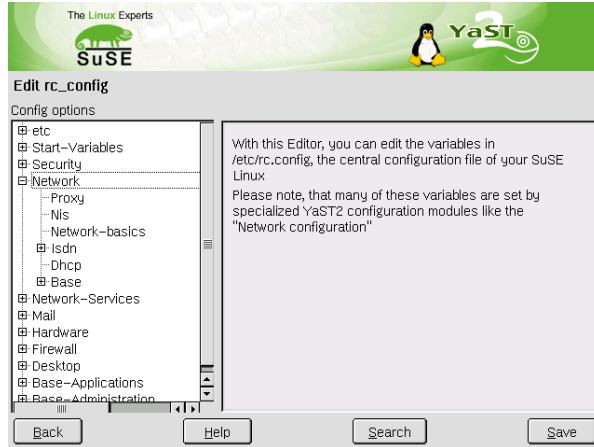



Figure 4.2: YaST2: Configuring the RC.Config Editor

‘Edit’ continues to an overview of all the services and daemons, supplemented with information as to whether they have been activated on your system and for which runlevels. Highlight a line with the mouse and activate the checkboxes for runlevels ‘0’, ‘1’, ‘2’, ‘3’, ‘5’, ‘6’, and ‘s’ and, with that, state which service or daemon should be activated for which runlevel. Runlevel 4 is undefined — this is always reserved for custom settings.

With ‘Start’ and ‘Stop’, decide whether a server should be implemented. The current status is checked via ‘Update’, if this has not already been done automatically. ‘Reset to default value’ allows you to restore the default settings to their initial state following installation. ‘Activate service’ will only appear if the service is currently disabled. ‘Reset all services to default value’ restores all services to their original state following installation. ‘Finish’ saves the system configuration.

Caution

Caution, no experiments here — this is an expert tool. If certain service settings are wrong, the system may not be able to be restarted, possibly requiring a reinstallation from scratch, or security gaps could turn up in your system.



4.6.3 Expert Partitioner

The partitioning module enables editing and deletion of existing partitions, as well as the creation of new ones. Access Soft RAID and LVM configuration from here. See page 94 and 89. In normal circumstances, partitions are specified during installation. However, it is possible to integrate a second hard disk in an already existing Linux system. First, the new hard disk must be partitioned.

Then it must be mounted and entered into the `fstab` file. It may be necessary to copy some of the data to move an `/opt` partition from the old hard disk to the new one.

Use caution if you want to repartition the hard disk you are working on — this is essentially possible, but you will have to reboot the system right afterwards. It is a bit safer to boot from CD then repartition it.

The ‘Expert options’ button reveals a pop-up menu containing the following commands:

Reset and Re-Read Rereads the partitioning from disk. For example, you will need this for manual partitioning in the text console.

Read old fstab This will only be relevant during installation. Reading the old `fstab` is useful for completely reinstalling your system rather than just updating it. In this case, it is not necessary to enter the mount points by hand.

Delete old Partition Table This completely overwrites the old partition table. For example, this can be helpful if you have problems with unconventional disk labels. Using this method, all data on the hard disk will be lost.

4.6.4 Logical Volume Manager (LVM)

The Logical Volume Manager (LVM) enables flexible distribution of hard disk space over several file systems. As it is difficult to modify partitions on a running system, LVM was developed: it provides a virtual “pool” (Volume Group — VG for short) of memory space, from which logical volumes (LV) can be generated if needed. The operating system will access these instead of the physical partitions.

Features:

- Several hard disks or partitions can be combined into a large logical partition.
- If a LV (e. g., `/usr`) is full, it can be enlarged with the appropriate configuration.
- With the LVM, even append hard disks or LVs in a running system. However, “hot-swappable” hardware, designed for these types of interventions, is required for this.

Implementing LVM already makes sense for heavily used home PCs or small servers. If you have a growing data stock, as in the case of databases, MP3 archives, or user directories, the Logical Volume Manager is just the right thing for you. This would allow you file systems that are larger than physical hard disk. Another advantage of the LVM is up to 256 LVs can be added. Keep in mind that working with the LVM is very different than working with conventional partitions.

Further information on configuring the “Logical Volume Manager” (LVM) can be found in the official LVM HOWTO at:

<http://www.sistina.com/lvm/Pages/howto.html>

or at: <http://www.suse.com/us/support/oracle/>.

LVM Configuration with YaST2

Activate the YaST2 LVM configuration by choosing ‘Custom Partitioning with LVM’ at the first step of preparing the hard disk before continuing with installation. Find the respective button in the YaST2 Control Center under ‘System’ on an installed system.

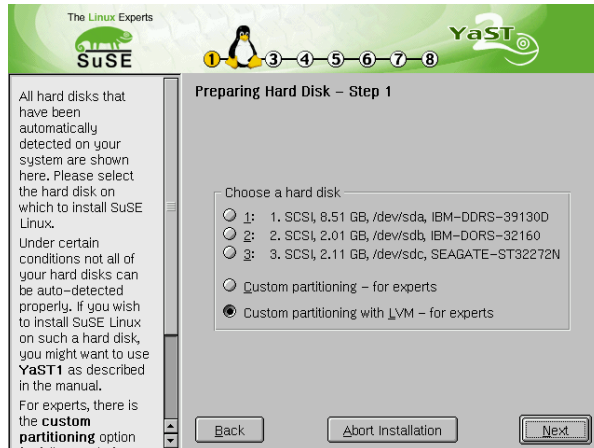


Figure 4.3: YaST2: Activating LVM During Installation

LVM — Partitioner

First, you will be presented with a dialog where you can repartition your hard disks. Delete or modify existing partitions here or add new ones. A partition to use for LVM must have the partition label 8E. These partitions are indicated by “Linux LVM” in the partition list.

You do not need to individually set the 8E label for all partitions designated for LVM. If need be, YaST2 will automatically set the partition label of a partition assigned to an LVM Volume Group to 8E. For any unpartitioned areas on your disks, create LVM partitions in this dialog. These partitions should then be designated the partition label 8E. They do not have to be formatted and no mount point can be entered in respect to these.

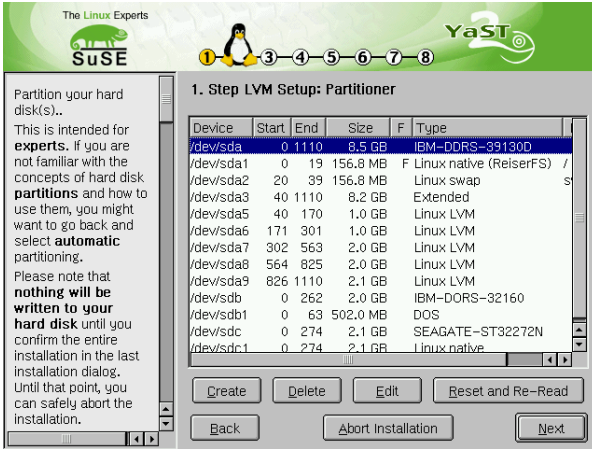


Figure 4.4: YaST2: LVM Partitioner

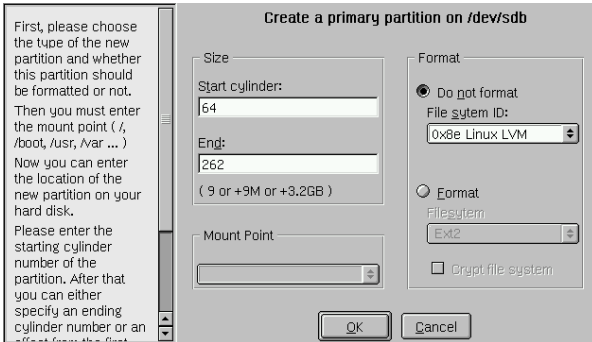


Figure 4.5: YaST2: Creating LVM Partitions

Note

If a working LVM configuration already exists on your system, it will be automatically activated as soon as you begin configuring the LVM. If this is successfully activated, any disks containing a partition belonging to an activated volume group can no longer be repartitioned. The Linux kernel will refuse to read the modified partitioning of a hard disk as long as only one partition on this disk is being used. Of course, repartitioning disks not belonging to an LVM volume group is not a problem at all. If you already have a functioning LVM configuration on your system, repartitioning is usually not necessary. In this screen, configure all mount points not located on LVM Logical Volumes. The root file system in YaST2, at least, must be stored on a normal partition. Select this partition from the list and specify this as root file system using the 'Edit' button. To ensure LVM's optimal flexibility, we recommend that you pool all additional file systems onto LVM logical volumes. After specifying the root partition, exit this dialog.

LVM — Configuring Physical Volumes

This dialog is responsible for managing LVM volume groups (often indicated by “VG”). If no volume group exists on your system yet, you will be prompted in a pop-up window to add one. “System” is suggested as a name for the volume group where the SuSE Linux system files are located. Physical Extent Size (often indicated by PE Size) defines the maximum size of a physical and logical volume in this volume group. This value is normally set to 4 megabytes. This allows for a maximum size of 256 gigabytes for physical and logical volumes. The physical extent size should only be increased if you need larger logical volumes than 256 gigabytes (e. g., to 8, 16, or 32 megabytes).

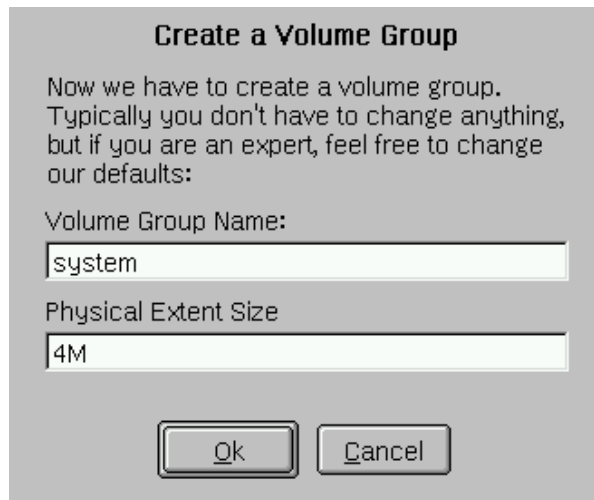


Figure 4.6: YaST2: Adding a Volume Group

In the following dialog, all partitions are listed which have either the "Linux LVM" or "Linux native" type. Therefore, all swap and DOS partitions will not be shown. If a partition is already assigned to a volume group, the name of the volume group will be shown in the list. Unassigned partitions are indicated by "_".

The volume group currently being edited can be modified in the selection box above, to the left. The buttons in the upper right enable creation of additional volume groups and deletion of existing volume groups. In any case, only volume groups can be deleted to which no other partitions are assigned. No more than one volume group needs to be created for a normally installed SuSE Linux system. A partition assigned to a volume group is also referred to as a physical volume (often indicated by PV). To add a previously unassigned partition to the selected volume group, first click on the partition then the 'Add Volume' button below the selection list. At this point, the name of the volume group is entered next to the selected partition. Assign all partitions reserved for LVM to a volume group. Otherwise, the space on the partition will remain unused. Before exiting the dialog, every volume group has to be assigned at least one physical volume.

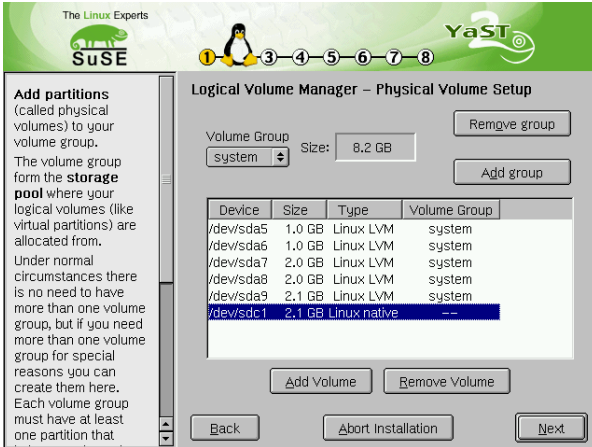


Figure 4.7: YaST2: Overview of the Partitions

Logical Volumes

This dialog is responsible for managing logical volumes (often indicated by just “LV”).

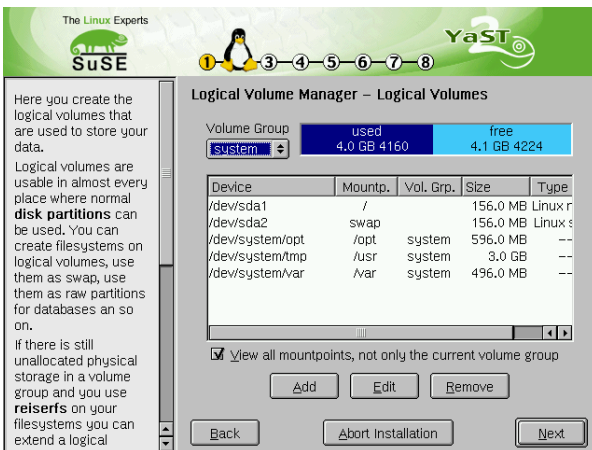


Figure 4.8: YaST2: Logical Volume Management

Logical volumes are assigned, one to each volume group, and have a particular given size. Normally, a file system is created on a logical volume (e. g., reiserfs, ext2) and is then designated a mount point. The files stored on this logical volume can be found at this mount point on the installed system. All normal Linux partitions to which a mount point is assigned, all swap partitions, and all already existing logical volumes are listed here. If you have already configured LVM on your system, the existing logical volumes will have to be entered now. Before you proceed, assign the appropriate mount point to these logical volumes.


If you are configuring LVM on a system for the first time, no logical volumes will be displayed in this screen yet. A logical volume must be generated for each mount point (using the ‘Add’ button). You will also need to specify the size, the file system type (e. g., reiserfs or ext2), and the mount point (e. g., /var, /usr, /home).

Figure 4.9: YaST2: Creating Logical Volumes

If you have created several volume groups, switch between the different volume groups in the selection list above and to the left. The added logical volumes are listed in the volume group displayed there. Once you have created all the logical volumes as required, the LVM configuration will be complete. At this point, exit the dialog and go on to software selection if you are still in the installation process.

Caution

Using LVM is associated with increased risk factors such as loss of data. Risks also include application crashes, power outage, and faulty commands. Save your data before implementing LVM or reconfiguring volumes — in other words, never work without a backup.



4.6.5 Soft RAID

The purpose of RAID (Redundant Array of Inexpensive Disks) is to combine several hard disk partitions into one large “virtual” hard disk for the optimization of performance and data security. Using this method, however, one advantage is sacrificed for another. “RAID level” defines the pool and common triggering

device of the all hard disks, known as the RAID controller. A RAID controller mostly uses the SCSI protocol, because it can drive more hard disks better than the IDE protocol. It is also better able to process parallel running commands.

Instead of a RAID controller, which can often be quite expensive, the Soft RAID is also able to take on these tasks. SuSE Linux offers the option of combining several hard disks into one Soft RAID system with the help of YaST2 — a very reasonable alternative to Hardware RAID.

Customary RAID Levels

RAID 0 This level improves the performance of your data access. Actually, this is not really a RAID, because it does not provide data backup, but the name “RAID 0” for this type of system has become the norm. With RAID 0, two hard disks are pooled together. The performance is very good — although the RAID system will be destroyed and your data lost, even if just one of the many remaining hard disks fails.

RAID 1 This level provides more than adequate backup for your data, since the data is copied to another hard disk 1:1. This is known as “hard disk mirroring” — if a disk is destroyed, a copy of its contents is located on another one. All of them except one could be damaged without endangering your data. The writing performance suffers a little in the copying process when using RAID 1 (ten to twenty percent slower), but read access is significantly faster in comparison to any one of the normal physical hard disks, because the data is duplicated so can be parallel scanned.

RAID 5 RAID 5 is an optimized compromise between the two other levels in terms of performance and redundancy. The hard disk potential equals the number of disks used minus one. The data is distributed over the hard disks as with RAID 0. “Parity blocks”, created on one of the partitions, are there for security reasons. They are linked to each other with XOR — thus enabling the contents, via XOR, to be reconstructed by the corresponding parity block in case of system failure. With RAID 5, no more than one hard disk can fail at the same time. If one is destroyed, it must be replaced as soon as possible to save the data.

Soft RAID Configuration with YaST2

Access Soft RAID configuration by way of the ‘RAID’ module under ‘System’ or via the partitioning module under ‘Hardware’.

1st Step: Partitioning

First, see a list of your partitions under ‘Expert Settings’ in the partitioning tool. If the Soft RAID partitions have already been set up, they will appear here. Otherwise, you will have to set them up from scratch. For RAID 0 and RAID 1, at least two partitions are needed — usually for RAID 1, exactly two and no more. If RAID 5 is being used, at least three partitions will be required. It is recommended to only take partitions of the same size. The RAID partitions

should be stored on various hard disks to insure against the risk of losing data if one is defective (RAID 1 and 5), as well as to optimize the performance of RAID 0.

2nd Step: Setting Up RAID

Click on 'RAID' and a dialog will appear where you can choose between RAID levels 0, 1 and 5. In the following screen, you will have the option of assigning the partition to the new RAID. The 'Expert options' button opens the settings options for the "chunk size" — fine-tune the performance. Clicking on the checkbox 'Persistent superbblock' ensures that the RAID partitions will be recognized as such when booting.

After completing the configuration, you will then see the `/dev/md0` device and others indicated on the expert page in the partitioning module by "RAID".

Troubleshooting

Find out whether a RAID partition has been destroyed by the file contents `/proc/mdstats`. The basic procedure in case of system failure is to shut down your Linux system and replace the defective hard disk with a new one partitioned the same way. Then restart your system and give the `raidhotadd /dev/mdX /dev/sdX` command. This will enable the hard disk to be integrated automatically into the RAID system and be fully reconstructed.

More details on this can be found in the HOWTOs at `/usr/share/doc/packages/raidtools/Software-RAID-HOWTO.html` and <http://www.LinuxDoc.org/HOWTO/Software-RAID-HOWTO.html> or over a Linux RAID mailing list, such as <http://www.mail-archive.com/linux-raid@vger.rutgers.edu>.

There you will find help if unexpected complications arise.

4.6.6 Kernel Selection

Caution

It is very risky to modify the kernel in an executable Linux system, as you can encounter compatibility problems, which may mean that your system may not even be able to be rebooted. Therefore, only experts should use this module.



4.7 Miscellaneous

'Miscellaneous' gives the option of making support requests, loading vendor driver CDs, and viewing the start protocol (`/var/log/boot.msg`) and the system protocol (`/var/log/messages`). Instructions for this can be found in the "Configuration" manual.

4.7.1 Printers for CUPS and LPD

The default setting for printing is the default print system which implements the BSD spooler LPD. package `lprold` is installed on your system for this purpose. With the help of YaST2 switch to the CUPS (Common Unix Printing System) printing system or from CUPS back to LPD. The old LPD configuration will be retained. Since their configuration files overlap somewhat, both printing systems are mutually exclusive — only one can be installed at a time. Follow the instructions in YaST2 for switching to CUPS and keep your CDs at hand.

More information on CUPS can be found at <http://www.cups.org/>



Note

The decisive factor for print quality is not whether you are using LPD or CUPS, but the related Ghostscript driver and the parameter settings.

5 Booting and Boot Managers

This chapter describes various methods of booting a Linux system. To have a better understanding of what is involved, we will first illustrate some technical details of booting a PC.

5.1 Booting a PC

After turning on your computer, the first thing that happens is that the BIOS (Basic Input Output System) takes control, initializes the screen and keyboard, and tests the main memory. Until this task is completed, no external devices or external storage media are known to the system.

Once the basic system has finished its internal setup, it starts to verify the hardware around it. Date, time, and information about some of the most important external devices are read from the CMOS settings (usually referred to as the CMOS setup). After reading the CMOS, the BIOS should recognize the first hard disk, including details such as its geometry. It can then start to load the operating system (OS) from there.

To load the OS, the system loads a 512-byte data segment from the first hard disk into main memory and executes the code stored at the beginning of this segment. The instructions contained there determine the rest of the boot process. This is the reason why the first 512 bytes of the hard disk are often called the *Master Boot Record* (MBR).

Although the whole process is rather complicated and this description is oversimplified, it should be clear that up, to this point (loading of the MBR), the boot sequence is independent of the installed operating system and is identical on all PCs. Also, all the PC has to access peripheral hardware is those routines (drivers) stored in the BIOS.

Master Boot Record

The layout of the MBR has been certified by an independent convention. The first 446 bytes are reserved for program code. The next 64 bytes offer space for a partition table for up to four partitions. See Section 2.5 page 36 and Section 2.6 page 39. Without the partition table, it would not be possible to create file systems (called drives under MS-DOS) on the hard disk — the disk would be virtually useless without the table. The last two bytes have to contain a special “magic number” (AA55). Any MBR which replaces this number by a different number is rejected.

Boot Sectors

Boot sectors are the first sectors on a hard disk partition, except for the extended partition which serves as a “container” for other partitions. They offer 512 bytes of space and are designed to contain code which is able to launch an operating system on this partition. Boot sectors of formatted DOS, Windows, and OS/2 partitions do exactly that. In contrast, Linux boot partitions are empty at the very start. A Linux partition cannot be started directly, although it may contain a kernel and a valid root file system. A valid boot sector follows the conventions and enters the “magic number” of the MBR into the last two bytes.

Booting DOS or Windows 95/98

The MBR contains information that determines which partition of a hard disk is “active” — which partition should be searched for the operating system to boot. Therefore, DOS has to be installed on the first hard disk. The executable code in the MBR (“first stage boot loader”) tests whether the marked partition contains a valid boot sector.

If this is the case, the “second stage boot loader” can be started from there. DOS system programs can now be loaded and you will see the usual DOS prompt. In DOS, only primary partitions can be marked active. Therefore, you cannot use logical partitions inside an extended partition as bootable DOS partitions.

5.2 Boot Concepts

The simplest boot concept affects only one machine with one operating system installed. A widely deployed PC configuration is DOS or Windows 95/98 as the only system installed. The boot sequence for this case has already been outlined. A similar concept can be used for Linux if it is the only operating system used. In this case, you could theoretically skip the installation of LILO. The big disadvantage of doing this is that you cannot pass additional parameters to the system kernel at boot time.

As soon as there is more than one operating system installed, there are a number of new boot possibilities.

Booting another OS from a floppy disk: One OS can be booted from the hard disk. Other operating systems can be booted by using boot disks.

- *Requirements:* the floppy drive must be bootable.
- *Example:* install DOS or Windows 95/98 on the hard disk and boot Linux from a floppy disk.
- *Advantage:* skip the potentially tricky boot loader installation.
- *Disadvantage:* you have to *ensure* that you are not running out of working boot disks.
- It might be an advantage or disadvantage that your Linux is not capable of booting without a boot disk, depending on for what it is used.

- The boot process will take slightly longer.

Boot chaining of additional systems: The same OS is always booted and others can optionally be started from within the first OS.

- *Requirements:* adequate programs for chain booting of operating systems must be available.
- *Example:* loading Linux from DOS using `loadlin` or starting a NetWare server from DOS with `server.exe`.

Installing a boot manager: Theoretically, this allows you to use an arbitrary number of operating systems on a single machine. The choice of systems is done at boot time. Changing operating systems requires a reboot.

- *Requirements:* the boot manager must work smoothly with all installed operating systems.
- *Examples* of coexisting boot managers (at least under certain circumstances) are OS/2 and the DOS boot loader `boot.sys`.

The following section describes the installation and configuration of a boot manager, using the Linux boot manager LILO. A complete description of LILO's features can be found in [Alm94].

This reference can be located and printed by entering

```
earth:/usr/share/doc/packages/lilo # lpr user.dvi
```

This is followed by a description of `loadlin`.

5.3 An Overview of LILO

The Linux boot loader is usually installed in the MBR (details below, page 103, and Section 5.5 page 110). When started, LILO already has access to both real mode hard disks and, due to its installation, is able to find all the data it needs from the *raw* hard drives without needing any information on partitioning. Because of this, operating systems can be booted from the first as well as from the second hard disk. The entries in the partition table that the standard DOS MBR uses to mark the active partition are ignored when using LILO in the MBR.

An important difference to the standard DOS boot sequence is that you can select any of the installed systems at boot time when using LILO. After loading the MBR into memory, LILO is started and you are asked to select one of the installed operating systems

What is LILO?

LILO is a versatile boot manager. It can launch an operating system in the following ways:

- by loading the boot sector of a partition and starting an operating system from this partition. This is what other boot managers also do.

- by loading the Linux kernel and starting Linux. This cannot be done by most other boot managers.

Furthermore, LILO provides an important option of being able to pass a command to the kernel. For security reasons, this can be protected totally or partially with a password.

Booting with LILO

When LILO is launched, it displays the text `LILLO` and a greeting message (which you entered during installation — see Section 5.4.2 page 107). Thereafter, a command prompt appears:

```
boot :
```

Here, select your operating system by entering its name. It is then booted. The name of the operating system has previously been set during installation. At this point, you can pass a parameter line to the Linux kernel. Get a list of all the operating system names available by pressing `(TAB)`.

The Components of LILO


Tip

The boot sectors installed by LILO contain a typical virus boot sequence. DOS virus scanners typically claim to have found the *AIRCOP* boot sector virus in files such as `/boot/any_b.b` or `/boot/any_d.b`. Also, you should disable any BIOS protection of the MBR.



The LILO machinery consists of the following components:

- the beginning, or *first step*, of the LILO code in a boot sector which activates the system boot
- the *heart* of the LILO code, localized in `/boot/boot.b`
- a *map* file, normally `/boot/map`, where LILO enters the location of Linux kernels and other data during its installation
- optional: a *message file*, the contents of which are displayed as a welcome message before the LILO boot selection. Its usual location is `/boot/message` (or similar)
- the different Linux kernel and boot sectors that LILO should offer



Caution

Any write access (even through file movements) on any of these files corrupts the map file, requiring a *reinstall* of LILO. This is only relevant when you change to another kernel.

Where to Install LILO

This is referring to the above-mentioned *first step* of LILO. Before going into detail, we would like to point out a very important general restriction: depending on the BIOS version on your computer, it may be required that the kernel image `/boot/vmlinuz` be located on the *first 1024 cylinders* of the hard disk. This can be achieved by a small extra partition which can be “mounted” in the directory `/boot`, all of which is located within the first 1024 cylinders.

This is because these are the only cylinders available when the BIOS starts the system. With an older BIOS and IDE drives, you can also be restricted to either of the first two hard disks (`/dev/hda` and `/dev/hdb`). If you have a very old BIOS, it is possible that the existence of (E)IDE hard disks could prevent your SCSI devices from booting.

Many newer BIOS versions allow access to additional devices, for example, in connection with EIDE hard disk controllers for up to 4 EIDE devices. Many modern SCSI controllers even allow SCSI devices to be “pushed to the front” to make them bootable. To use this feature with LILO, have a look at the `disk` options on page 108.

For the sake of simplicity, we will combine all this under the heading “1024-cylinder limit”. It should all be considered *before* a first-time installation (Section 2.4.1 page 28) — afterwards it may be too late and may cause a lot of extra work. For more information, see Section 5.8.2 page 119.

The following locations are acceptable for storing the LILO *boot sector*.

- **on a floppy disk.** This is the most secure, but also the slowest alternative for booting with LILO (see Section 5.6 page 113). Choose this alternative if you do not want to change boot sectors.
- **in the boot sector of a primary Linux partition on the first hard disk.** This leaves the MBR untouched. Before it can be booted, the partition has to be marked active, with `fdisk`. If Linux is fully installed on logical drives or partitions on the second hard disk, there is only the boot sector of the extended drive of the first drive left, if there is one. Linux `fdisk` can also activate such a partition.

If you want to boot multiple systems from the hard disk, this is quite long-winded. Every time you want to boot, you have to activate the corresponding boot sector *beforehand*. The next two variants are much less cumbersome.

- **in the *Master Boot Record*.** This variation offers the highest flexibility. Moreover, this is the only alternative possible if all of the Linux partitions reside on the second hard disk and there is no extended partition on the first drive. Every setting of the MBR must be edited with extreme care since errors may have severe consequences. The safety aspects are described in Section 5.5 page 110.
- **If you have used another boot manager until now** and you want to continue using it. Depending on its flexibility and power, there are several variations. A common case: you have a primary Linux partition on the second hard disk where you boot Linux. Your boot manager is able to boot this partition via a boot sector. Then you can activate your Linux partition by installing LILO into this boot sector and telling your boot manager that it is active.

Caution

Be careful if you try to make a *logical* Linux partition bootable by installing LILO onto it. Success is *not guaranteed*, even if your other boot manager is able to launch logical partitions.



Try it if you like. The safest way is to try it with a tiny Linux installation to see if it will work. Perhaps you will be lucky. The recommended way is still to create a primary and bootable Linux partition.

5.4 Configuring LILO

LILO is a flexible boot manager that offers many ways of adapting a configuration to your needs. The most important options and meanings are described below. If you want more detail, look at [Alm94].

Configuration of LILO is done in `/etc/lilo.conf`. If you are installing LILO for the first time, use YaST to configure LILO. You can fine-tune by editing `/etc/lilo.conf` at a later stage.

Note

`/etc/lilo.conf` should only be readable for `'root'`, as it could contain passwords (see Section 5.4.2 page 107). This is the default setting with SuSE Linux. If in doubt, just check by invoking the following command as root:

```
earth: # chmod 0600 /etc/lilo.conf
```



It is recommended to keep any existing old (and working) `lilo.conf` backed up in a safe place. Your settings only take effect when you reinstall LILO after changing `/etc/lilo.conf` (see Section 5.5 page 110).

5.4.1 Structure of `lilo.conf`

`/etc/lilo.conf` starts with a **global section** followed by one or more **system sections** for each operating system LILO should start. A new section is started by a line beginning with either **image** or **other**.

The order of entries in `/etc/lilo.conf` only matters in as much as the first one in the list is booted by default if no user interaction is taken — this can be set to **delay** and **timeout**.

A sample configuration for a machine with both DOS and Linux is shown in File 5.4.1. There are two Linux kernels (`/boot/vmlinuz` and `/boot/vmlinuz.old`) on `/dev/hda7`, as well as MS-DOS (or Windows 95/98) on `/dev/hda1`. The program MemTest86 is also available.

```
### LILO global section
boot      = /dev/hda          # LILO installation target: MBR
backup    = /boot/MBR.hda.990428 # backup file for the old MBR
                                     # 1999-04-28
vga       = normal          # normal text mode (80x25 chars)
read-only
menu-scheme = Wg:kw:Wg:Wg
lba32                    # Use BIOS to ignore
                                     # 1024 cylinder limit

prompt
password = q99iwr4        # LILO password (example)
timeout  = 80             # Wait at prompt for 8 s before
                                     # default is booted
message  = /boot/message  # LILO's greeting

### LILO Linux section (default)
image    = /boot/vmlinuz   # Default
label    = linux
root     = /dev/hda7       # Root partition for the kernel
initrd   = /boot/initrd

### LILO Linux section (fallback)
image    = /boot/vmlinuz.suse
label    = suse
root     = /dev/hda7
initrd   = /boot/initrd.suse
optional

### LILO other system section (DOS/Windows)
other    = /dev/hda1      # Windows partition
label    = windows

### LILO memtest section (memtest)
image    = /boot/memtest.bin
label    = memtest86
```

File 5.4.1: Sample Configuration of `/etc/lilo.conf`

Anything between a '#' and the end of a line is regarded as a comment. Spaces and comments are ignored by LILO and can be used to improve readability.

The most important lines, step-by-step, are:

- **Global section** (Parameter part)

- **boot=<bootdevice>**

The device on whose first sector LILO should be installed.

<bootdevice> may be: a floppy disk drive (`/dev/fd0`), a partition (e.g., `/dev/hdb3`), or an entire disk (e.g., `/dev/hda`). The last means installing LILO in the MBR. Default: if this option is missing, LILO is installed on the current root partition.

- **lba32**

With this option, ignore the 1024-cylinder limit of LILO if your BIOS supports this.

- **prompt**

Forces the LILO prompt to be displayed. The default is: no prompt (compare with **delay** further down). This is recommended if LILO needs to manage more than one system. In addition, **timeout** should be set to guarantee an automatic reboot if nothing is entered at the prompt.

- **timeout=<tenth-seconds>**

Sets a timeout for the prompted option, enabling an automatic reboot if no entry occurs in the given time. <tenth-seconds> is the remaining time in 0.1 s. increments. Pressing (**Shift** ↑) starts the timeout over. Default: infinite — no automatic reboot.

- **Linux section**

- **image=<kernelimage>**

Enter the name of the kernel image to be booted, including its directory location. With your new system, this is most likely `/boot/vmlinuz`, `/vmlinuz` for older SuSE Linux systems.

- **label=<name>**

This name has to be unique in `/etc/lilo.conf`. Otherwise, you can freely choose a name for the system (e.g., **Linux**). Maximum length is 15 characters. Use only letters, numbers, and underscore for names — no blanks or special characters, for example. For more on the specific rules for which characters to use, see [Alm94], as described in Section 5.2 page 101. The default is the file name of the kernel image (e.g., `/boot/vmlinuz`).

By entering this name at the LILO prompt, select which system to boot. It is recommended, if there are many systems installed, to keep track of them in a special message file (**message**).

- **root=<rootdevice>**

This is to give the kernel the name of the root partition (e.g., `/dev/hda2`) of your Linux system. This is recommended for security reasons. If this option is omitted, the kernel takes its own root partition <kernelimage>.

- **Linux part (Fallback)**

Even if you installed a customized kernel, you are still able to boot the SuSE standard kernel to start a stable system.

- **optional**

If you decide to delete `/boot/vmlinuz.suse` (*not recommended*), this section will be skipped without an error message during LILO installation.

- **Other systems**

- **other=<partition>**

other tells LILO to start the partitions of other systems (e.g., `/dev/hda1`).

- **label=<name>**

Name (your own choice) for the system. Recommended, because the default — the raw device name — is less informative.

5.4.2 Other LILO Configuration Options

The previous section covered the entries required in `/etc/lilo.conf`. Other useful options are discussed below.

Those options that are marked as image options belong to the appropriate section of the operating system. The others are intended for the global parameter section of `/etc/lilo.conf`.


- **backup=<backup>**

The file where LILO backs up the boot sector. The default is `/boot/boot.xxxx`, where `xxxx` is the internal device number of the installation partition.

We do not recommend use of a cryptic name (see our example above). You will not be able to use the implemented uninstall feature of LILO, but we think it is better to do this carefully by hand, anyway. (see Section 5.5 page 111)

Caution

If the backup file exists, LILO does *not* create a new one. Make sure you use a name not already in use.



- **compact**

This option is recommended if you want to install LILO onto a floppy disk. If enabled, LILO tries to read more sectors at a time, resulting in a faster boot process. This does not work on every machine. We do not recommend that you set this as the normal way is safer and it only provides a difference of one or two seconds.

- **loader=<boot loader>**

To load a boot sector that belongs to another operating system, LILO constructs a *pseudo MBR* in its map file. At boot time, LILO first starts this pseudo MBR, which, in turn, starts the other boot sector. This option specifies the file where the code for the pseudo MBR is found.

Default: /boot/chain.b (usually, this is correct).

Sometimes, another OS that needs to be booted from the first hard disk (e. g., DOS) is supposed to boot from another hard disk using LILO. There are additional options that cause the hard disks to swap according to their device numbers:

map-drive=<number> and **to=<number>**. See File 5.4.2.

The loader `os2_d.b` serves to load OS/2 from the second hard disk. ¹ *New in LILO Version 20:* “switching” devices has to be set explicitly (see File 5.4.2)

```
# Booting DOS from the second hard disk
# DOS bootable partition config begins
other = /dev/hdb1
  label = DOS
  loader = /boot/chain.b
    map-drive = 0x80 # first hd: BIOS number 0x80
    to         = 0x81 # second hd: BIOS number 0x81
    map-drive = 0x81
    to         = 0x80
  table = /dev/hdb
# DOS bootable partition config ends
```

File 5.4.2: /etc/lilo.conf Extract: Booting DOS from Second Hard Disk

- **table=<ptable>**

<ptable> sets the source device for the partition table written into the pseudo MBR (normally /dev/hda or /dev/sda).

- **disk=<device file>**

bios=<BIOS device number>

cylinders=<amount>

heads=<amount>

sectors=<amount> Tell LILO precisely which BIOS device number and geometry it should use. This is rarely needed. There is one major exception: *IDE-SCSI* system. If you own a BIOS that is capable of switching the boot devices *SCSI prior to IDE* and you want to use this feature, tell LILO the switched order from the perspective of the BIOS. This is achieved by an extra entry in the global section of `lilo.conf`. An example for a system with one SCSI and one IDE disk can be seen in File 5.4.3 on the next page.

¹ `any_b.b` (Booting from B:) and `any_d.b` (Booting from second hard drive) are obsolete from LILO Version 20.


```
# Enable LILO to correctly access /dev/sda and /dev/hda
# at boot time if their boot order is interchanged in
# the BIOS:
disk = /dev/sda      # The SCSI disk is regarded as ...
    bios = 0x80      # ... first BIOS disk;
disk = /dev/hda      # the IDE disk is regarded as ...
    bios = 0x81      # ... second BIOS disk.
```

File 5.4.3: `lilo.conf` Extract: Boot Order: SCSI Before IDE

- **linear**

Giving this option when installing LILO causes all references to hard disk sectors to be stored as logical instead of physical addresses, so that they are independent of any hard disk geometry. This option is intended for cases where, when booting, the BIOS detects a different geometry than that of the Linux system running. Only needed in rare cases.

The **linear** option does *not* release you from the constraints of the 1024-cylinder limit, which is determined by the BIOS geometry of the boot hard disk. Refer also to /usr/share/doc/sdb/de/html/kgw_lilo_linear.html.

- **message=<message-file>**

Points to a text file that should be shown on screen at system bootup. The file should not be longer than 24 lines (to prevent the message from scrolling) and may contain e. g. an overview of the LILO boot selection. This is recommended.

SuSE Linux normally uses a PCX image instead of the classical start-up message. To learn more about this, read file:/usr/share/doc/sdb/de/html/jkoeke_bootgrafik.html.

Note

If this option is set, the message file is then part of the LILO boot machinery and, after every change to this file, LILO has to be reinstalled (Section 5.5 on the following page).



- **password=<password>**

May be located either in a global or system-specific section. Provides secure access to LILO services or booting the corresponding system by means of a password. If you take this seriously, remove the password from `lilo.conf` after you have used it for the first time. As `'root'`, you can set a new password for LILO any time you like (you just need to reinstall it afterwards). It is recommended to also set the option **restricted**, otherwise it could be possible to launch a shell, see the man page about `lilo.conf` (**man lilo.conf**).

- **read-only**

This option tells the kernel to initially mount the root partition read-only, normal when starting Linux systems. If this is omitted, the kernel uses its internal settings.

- **delay=<tenth-seconds>**

If the prompt is *not* explicitly set, order a prompt by pressing **Shift** ↑, **Ctrl**, or **Alt**. The **delay=** option sets the time to elapse before LILO boots the first system in its list. The default is 0 — no waiting.

The **delay** option has no effect if a prompt is specifically requested by **prompt**.

- **vga=<mode>**

Selects VGA mode at start-up. Valid modes are **normal** (80x25), **ext** (80x50), or **ask** (asks when booting). For a framebuffer-enabled kernel, possible values are listed and described in `/usr/src/linux/Documentation/fb/vesafb.txt`.

- **append="<parameter>"** Image option for Linux kernel. Enables kernel parameters to be specified, for example, for hardware components, in the same way that this is possible at the LILO prompt. The kernel first reads the **append** line, then the prompt. Therefore, if the parameters given at the **append** line are different from those specified at the prompt, the latter ones will be used.

5.5 Installing and Uninstalling LILO

During a new Linux installation or at a later time, YaST will lead you through the steps of installing LILO interactively. In this section, we assume that some action is required that goes beyond what YaST can accomplish and we take a closer look at how LILO works during the installing and uninstalling process.

Caution

The installation of a boot manager is tricky. Ensure *in advance* that you are completely able to boot Linux and other mounted systems. You must have **fdisk** installed on a crash recovery disk, otherwise you might find yourself in the awkward situation of not being able to access your hard disk at all.



Installation After Changing the Configuration

If any of the LILO components have changed or you have modified your configuration in `/etc/lilo.conf`, reinstall LILO. This is easily done by launching the *Map Installer* like this:

```
earth: # /sbin/lilo
```

What happens now is that LILO writes a backup of the target boot sector, writes its *first step* into it, and creates a new map file (see also Section 5.3 page 102). LILO now announces each installed system — for an example, see Output 5.5.1.

```
Added linux *
Added suse
Added windows
Added memtest86
```

Output 5.5.1: Output After Launching LILO

When the installation is complete, the machine can be rebooted:

```
earth: # shutdown -r now
```

During reboot, the BIOS first performs its system test. Directly afterwards, you will see LILO and its command prompt, where you can enter parameters and select a boot image from the recently installed configurations. (TAB) shows a list of all systems installed.

Installation After Recompiling a Kernel

If you want to include a freshly created kernel into your LILO boot setup, the Linux kernel Makefile offers an all-in-one solution. All the commands to configure and create the kernel are put together in the file `/usr/src/linux/Makefile`. Here, the `INSTALL_PATH=/boot` is specified (see Section 10.5 page 233). This Makefile has a **target** called `bzlilo` which, after a kernel compilation, automatically copies the currently installed kernel `/boot/vmlinuz` (this used to be `/vmlinuz`) to `/boot/vmlinuz.old`, the new kernel to `/boot/vmlinuz`, then reinstalls LILO. This can be done by entering the command:

```
earth:/usr/src/linux # make bzlilo
```


instead of `make zimage`. This is only useful if you have edited `/etc/lilo.conf` *in advance* and if your current kernel really is located in `/boot/vmlinuz`. The new, as well as the old, kernel should now be listed. See File 5.4.1 page 105 for an example of the resulting `/etc/lilo.conf`.

At the LILO prompt, launch either of the two kernels. This makes your boot more secure, because you can still boot your old kernel even if the new one fails. For more on creating a new kernel, see Chapter 10 page 229.

Uninstalling LILO

Caution

Uninstalling a boot manager is tricky. Ensure *in advance* that you are completely able to boot Linux and other systems with their respective boot disks. You should have `fdisk` installed on every boot disk, otherwise you might find yourself in the unfortunate situation of not being able to access your hard disk at all.



Perhaps one day it will be necessary for you to uninstall LILO. This is accomplished by writing back the target boot sector where LILO has been installed. This is not a problem in Linux *if* there is a valid backup (see Section 5.4.2 page 107, option **backup**).



Caution

A boot sector backup is no longer valid if the partition in question has a new file system (for DOS users: has been formatted). The partition table of an MBR backup becomes invalid if the hard disk in question has been repartitioned in the meantime. Obsolete “backups” are time bombs. It is best to delete them as soon as possible. Unpacking old and invalid backups into system sectors is a direct route to data loss!

It is very simple to get back a DOS, Windows 95/98, or OS/2 MBR. Just enter the MS-DOS command (available since 5.0)

```
C:\> FDISK /MBR
```

or, on OS/2,

```
C:\> FDISK /NEWMBR
```

These commands only write the first 446 bytes (the boot code) into the MBR and leave partitions untouched.

For other restorations, first make a backup of the LILO sector in question — just to be on the safe side. Now you should check, at least twice, whether your old backup file is the correct one and if it is exactly 512 bytes in size. Finally, write it back, but do not confuse **if=** and **of=**.

- If LILO resides in partition *yyyy* (e. g., *hda1*, *hda2*,...):

```
earth: # dd if=/dev/yyyy of=New-File bs=512 count=1  
earth: # dd if=Backup-Date of=/dev/yyyy
```

- If LILO resides in the MBR of *zzz* (e. g., *hda*, *sda*):

```
earth: # dd if=/dev/zzz of=New-File bs=512 count=1  
earth: # dd if=Backup-Date of=/dev/zzz bs=446 count=1
```

The last command is “cautious” and does not overwrite the partition table. Again, *do not forget*: with **fdisk**, mark the desired starting partition as *bootable*. By the way, note how easy and fast a boot sector backup is done. We recommend you do this frequently.

5.6 Creating a Linux Boot Disk

A Linux boot disk consists (somewhat simplified) of one or more Linux kernels, possibly managed by LILO. It serves to start your system even if it is not possible to boot directly from hard disk (possible reasons: overwritten MBR, misconfigured boot manager, errors while installing).

A boot disk such as this loads *only* the kernel. Everything else, including working system programs and `init` start scripts, must be provided by the installation on the hard disk. The connection between the boot disk and the system on the hard disk is established by the fact that, in the kernel, the root partition in question is set as the root device.

Do not confuse this with the SuSE boot disk used for installation and emergencies. To create a new SuSE boot disk, copy the appropriate image from the directory `disks` on the SuSE CD-ROM to a floppy disk (see Section 13.6 page 297).

Boot Disk Without LILO

Normally, if your hard disk controller support has been integrated into your kernel, your kernel does *not* need any additional command line with hardware information. The easiest way to create a boot disk is to just write the actual kernel onto a raw disk and adjust the root device, if this has not been done before.

```
earth: # /sbin/badblocks -v /dev/fd0 1440
earth: # dd if=Your_Kernel of=/dev/fd0 bs=18k
earth: # rdev /dev/fd0 Your_Root_Partition
earth: # rdev -R /dev/fd0 1
```

The last command makes sure that the kernel initially mounts root as read-only (the start-up scripts expect this).

Boot Disk with LILO

You can create a much more capable boot disk with a greeting, prompt, kernel parameters, and other LILO goodies by transferring the complete LILO booting start machinery onto the disk (see Section 5.3 page 102). For this, the disk needs a file system. The Minix file system is best suited for this.

To do this, proceed as follows:

- Create a Minix file system on a new and empty floppy disk and mount the disk to, for example, `/mnt` with the commands:

```
earth: # /sbin/mkfs.minix -c /dev/fd0 1440
earth: # /bin/mount /dev/fd0 /mnt
```

- Now copy your kernel files and the LILO file `/boot/boot.b` to `/mnt`, for example, onto the floppy disk.
- Optional: create a message file `/mnt/message`.
- Create `lilo.conf` on `/mnt`. Adapt this to your needs (give the correct name of the kernel, etc.). See File 5.6.1 on the next page for an example.
- If you are using an `initrd`, add it to `lilo.conf` as described in Section 13.4.3 page 289.
- Install LILO with *this* `lilo.conf`:

```
earth: # /sbin/lilo -C /mnt/lilo.conf
```

```
# LILO Configuration file bootdisk
# Start LILO global Section
boot=/dev/fd0          # Installation: Floppy
install=/mnt/boot.b   # Of course LILO and
map=/mnt/map          # map file onto floppy!
message=/mnt/message  # optional
prompt
timeout=100          # Wait at prompt: 10 s
vga = normal         #
# End LILO global section
#
# Linux bootable partition config begins
image = /mnt/vmlinuz  # default
    root = /dev/Your_Root_Device # Here is your root partition!
    label = linux
# Linux bootable partition config ends
#
# System sections for more kernel here:
```

File 5.6.1: `lilo.conf` for a Boot Disk

- Unmount the floppy
earth: # `/bin/umount /mnt`
- Do not forget to check your boot disk at the next system start to check whether it works or not.

5.7 Sample Configurations

If Linux is the only operating system on your machine, there is nothing to do, since everything needed has already been done by YaST. The following sections give some example configurations for multisystem computers. Further information can be found under `/usr/share/doc/howto/en/mini/Linux+*.gz`.

5.7.1 DOS/Windows 95/98 and Linux

Requirements: There must be at least a primary partition each for DOS/Windows 95/98 and Linux which is below the 1024--cylinder limit (Section 5.3 page 103). For this case, we have already discussed a configuration (File 5.4.1 page 105) — only the settings for **root**, **image**, and **other** have to be adapted. LILO is installed in the MBR.

Save your `/etc/lilo.conf` and be sure you have a Linux boot disk. Windows 95/98 is especially inclined to eliminate “foreign” MBRs. If you can still boot Linux using your boot disk, this problem is quickly solved with the command

```
earth: # /sbin/lilo
```

which will complete your LILO installation.

5.7.2 Windows NT and Linux on One Hard Disk

1. If Windows NT and Linux need to coexist on the same hard disk, use the NT boot manager for booting. This can either start the kernel images or the boot sectors themselves. Execution of the following steps prepares everything for a peaceful coexistence of Linux and Windows NT:
 - Install NT.
 - Partition the NT disks (using FAT so that Linux can write on it).
 - Install Linux as usual (in our example, the root partition is on `/dev/sda3`). Mount either the DOS partition or an error-free DOS floppy disk (for example, on `/dos`).
 - Install LILO, but install it in Linux's root partition (`/dev/sda3`), **not** in the MBR (`/dev/sda`). You may still configure a selection of Linux kernels for LILO. See File 5.7.1 for an example `lilo.conf`.

```
# LILO Configuration file
# Start LILO global Section
boot=/dev/sda3           # Target of installation
backup=/boot/boot.sda3.970428 # Backup previous boot sector;
                           # 28. Apr 1997

prompt
timeout=100             # Wait at prompt: 10 s
vga = normal            # force sane video state
# End LILO global section
# Linux bootable partition config begins
image = /vmlinuz        # default image to boot
  root = /dev/sda3      # Here the root partition!
  label = Linux
# Linux bootable partition config ends
```

File 5.7.1: `lilo.conf` for Booting a Linux Root Partition

- Copy the LILO boot sector to a location where NT can find it. For example,


```
earth: # /bin/dd if=/dev/sda3 bs=512 count=1 of=/dos/bootsek.lin
```

This step, as well as the following, has to be performed after every kernel update.

- Boot NT. Copy `bootsek.lin` from data disk to main directory of NT's system drive, if it is not already there.
- In `boot.ini` (first setting attributes), supplement at the end:

```
c:\bootsek.lin="Linux"
```

- After the next boot (if everything went smoothly), there should be an entry in NT's boot manager.
2. Another possibility: install LILO in the MBR and claim that it is DOS to Windows NT (as in our previous example).



Caution

NT 3.5* and 4.0 do not recognize Linux partition types 82 and 83. Make sure that no NT program tries to “repair” your partition table. This would result in loss of data! Always have valid backups of the LILO boot sector at hand.

5.7.3 OS/2 and Linux

1. Use the *OS/2 boot manager* for booting. It can launch unlimited primary and logical partitions provided they are below the 1024--cylinder limit. The user is responsible for the partitions. This boot manager is configured by OS/2's `fdisk`.

Preparation on the Linux side: consists merely of making a partition bootable (usually this is the root partition) with LILO. You can use the same `lilo.conf` as in the Windows NT example, but there is one thing that you have to consider *in advance*...

Preparation on the OS/2 side: OS/2 not only uses the conventional and obvious entries for existing partitions in MBRs on hard disks, but also uses “waste” space in these sectors for additional information. If these are inconsistent, OS/2's `fdisk` considers these partitions faulty and refuses to provide boot manager services. The `fdisk` commands of other systems do not know how to handle these extra partitions. Conflicts are inevitable.

Therefore: before installing Linux, load OS/2 (the installation system is sufficient) and create the Linux partitions with OS/2's `fdisk`, at least the logical drives. This initially creates additional OS/2 partitions, which might get in the way.

Solution: immediately after creating the partitions, load the Linux installation system (or the rescue disk from the SuSE Linux CD) and change the partition types to 83 (Linux native) using `fdisk`. Now these partitions will be ignored by OS/2.

2. *Second option:* LILO is used as the main boot manager on a primary partition on the first hard disk. This is a special case in the following example, which also includes DOS. Using the MBR is less advantageous, because repartitioning with another `fdisk` could erase the MBR and thus remove LILO. This special case is also considered in our next example, where DOS is additionally involved.

5.7.4 DOS, OS/2, and Linux

1. If you have used the *OS/2 boot manager* for DOS and OS/2 and want to continue using it, simply add Linux to its start menu as described in the previous example.

2. Another possibility is to use *LILLO* as the main boot manager on a primary partition of the first hard disk. The following complicated example of `lilo.conf` (File 5.7.2) assumes that the DOS boot partition (primary) and the Linux boot partition (also primary) are both on the first hard drive — each of them within the 1024-cylinder limit of the two partitions. OS/2 is installed on a second hard disk. This is why a special loader (`/boot/os2_b.b` is used instead of `/boot/chain.b`).

It does not matter whether the MBR code is DOS or OS/2. The important thing is that the LILO boot partition (`/dev/sda4`) must be marked as active in the partition table with `fdisk`.

```
# LILO Configuration file
# Start LILO global Section
boot = /dev/sda4          # LILO in Linux root partition
backup = /boot/boot.sda4.970428
message = /boot/message # Greeting message
prompt
delay = 100
vga = normal
#
# Linux bootable partition config begins
image = /vmlinuz
    label = linux
    root = /dev/sda4
# Linux bootable partition config ends
#
# OS/2 bootable partition config begins
other = /dev/sdb5
    table = /dev/sdb
    label = os2
    loader = /boot/os2_b.b
# OS/2 bootable partition config ends
#
# DOS bootable partition config begins
other = /dev/sda1
    table = /dev/sda
    label = dos
# DOS bootable partition config ends
```

File 5.7.2: LILO with DOS, OS/2, and Linux on Two Hard Disks

5.8 LILO Problems

Some Guidelines

Some simple guidelines at the beginning will avoid most LILO problems in advance (this is taken from the LILO documentation [Alm94]):

- *Don't panic!* If anything does not work, try to find the error or the cause first. Check the diagnosis before you start fixing the problem.
- Always have an up-to-date and tested *boot disk* at hand.
- SuSE Linux contains a full Linux system on its boot disk and installation CD (for the rescue system, see Section 13.6 page 297) to allow you to reach all your Linux partitions. Tools are included for repairing almost any problems that can occur.
- Read the complete LILO documentation, especially if the system does not do what you want it to do.
- Check `/etc/lilo.conf` *before* using the map installer (`/sbin/lilo`).
- Be careful if you are using a large hard disk or multiple ones. Be aware of the 1024-cylinder limit.
- Try with and without the `linear` option (normally it should be better without).

5.8.1 Diagnosis of Errors: LILO Start Messages

This is mainly Section 5.2.1 from [Alm94].

When LILO loads itself, it displays the word `'LILO'`. Each letter is printed before or after performing some specific action. If LILO fails at some point, the letters printed so far can be used to identify the problem.

nothing No part of LILO has been loaded. Either LILO is not installed at all or the partition on which its boot sector is located is not active.

'L' ***error*** ... The *first stage* boot loader has been loaded and started, but it cannot load the second stage boot loader (`/boot/boot.b`). The two-digit error codes indicate the type of problem. This condition usually indicates a media failure or a geometry mismatch.

'LI' The second stage has been invoked, but could not be started. This can either be caused by a geometry mismatch or by moving `/boot/boot.b` without reinstalling LILO.

'LIL' The second stage of boot loader has been started, but it cannot load the descriptor table from the map file. This is typically due to a physical error of the boot device or a faulty disk geometry.

'LIL?' The second stage boot loader has been loaded at an incorrect address. This is typically caused by a subtle geometry mismatch or by moving `/boot/boot.b` without reinstalling LILO.

'LIL-' The descriptor table in the map file is corrupt. This can either be caused by a geometry mismatch or by moving `/boot/boot.b` without reinstalling LILO.

'LILO' All parts of LILO have been successfully loaded.

The most common causes for *geometry errors* are not physical defects or invalid partition tables, but errors in LILO installation, including:

- disregarding the 1024--cylinder limit (see next section)
- an unsuccessful attempt at starting LILO from a logical partition


5.8.2 The 1024--Cylinder Limit



Note

Recently, BIOS versions are available which enable you to start operating systems above the 1024--cylinder limit. The current LILO version can use this BIOS extension. YaST and YaST2 will inform you accordingly of these options for your BIOS while configuring LILO. If your BIOS does not include this extension, continue reading here.

As emphasized before (on page 103), the entire LILO machinery (including all data needed for booting) must be able to process BIOS calls, which means it must reside below the 1024--cylinder limit on the hard disk. The sections of the hard disk that can be used, called *allowed sections*, have already been discussed. This restriction affects *only* the boot-up machinery. It is not required that LILO be installed on the Linux root partition. It is even possible, but quite dangerous, to put the boot machinery onto partitions of other operating systems to which Linux has read and write access.



Caution

Never install the LILO boot sector onto an unknown partition because you will severely damage the file system

The best method is to create a primary partition (within the allowed section) and to install all LILO files (including the LILO boot sector) into this partition. This will be, in most cases, the Linux root partition.

You can also add it to `/boot` with YaST. The only condition is that there has to be enough space for `boot.b`, `map`, `message`, and the Linux kernels that LILO should boot.

A few megabytes is enough. It does not matter where you put the rest of your partitions. There are no more restrictions. As soon as the kernel runs, you have unrestricted access to all installed drives.

But what to do if there is no space for such a partition? If you neither want to repartition your hard disk, upgrade to SCSI, or purchase a new BIOS version, there are still two (makeshift) possibilities:

- Use a boot disk instead of LILO on the hard disk or, if you are also running MS-DOS, use `loadlin`.

- Install the LILO boot machinery onto a Linux partition in the permitted section and where Linux has write access (e. g., a FAT or VFAT drive). We cannot put the LILO boot sector there as well. So there are only two places to put it. Either at the start of an extended partition on the first drive — as long as it is beneath the 1024--cylinder limit — or on the MBR.

Suppose that the partition in question is mounted on `/mnt`, that LILO is installed in the MBR (`/dev/hda`), and that you also boot DOS from `/dev/hda1`. Proceed as follows:

- Create a new directory (e. g., `/mnt/LINUX`) and copy the LILO files mentioned above to it (`boot.b`, `map`, `message`) as well as the chain loader of other operating systems (normally `chain.b`) and the Linux kernels that LILO should boot.
- Create a `/mnt/LINUX/lilo.conf` where all paths point to `/mnt/LINUX` (see File 5.8.1).

```
# LILO Configuration file
# Start LILO global Section
boot=/dev/hda          # Installation target
backup=/mnt/LINUX/hda.xxxx # backup of old MBR
install=/mnt/LINUX/boot.b # Of course LILO and
map=/mnt/LINUX/map     # map file are in /mnt/LINUX!
message=/mnt/LINUX/message # optional
prompt
timeout=100          # Wait at prompt: 10 s
vga = normal        #
# End LILO global section
#
# Linux bootable partition config begins
image = /mnt/LINUX/First_Kernel # default
    root = /dev/Your_Root_Device # Root partition!
    label = linux
# Linux bootable partition config ends
#
# System section for other kernels:
#
# End Linux
# DOS bootable partition config begins
other = /dev/hda1     # MSDOS system drive
    label = dos
    loader = /mnt/LINUX/chain.b
    table = /dev/hda
# DOS bootable partition config ends
```

File 5.8.1: `lilo.conf` for Other Partitions

- Install LILO with *this* `lilo.conf`:
`earth: # /sbin/lilo -C /mnt/LINUX/lilo.conf`
After that, LILO should work. Boot MS-DOS and protect the LILO files as well as possible against write access (any write access disables LILO).

To accomplish this, assign to all files in `X:\LINUX` (where the `X` is the DOS drive mounted to `/mnt`) the DOS attributes *system* and *hide*.

In conclusion, we point you toward two HOWTOs in `/usr/share/doc/howto/en/mini/` — `LILO.gz` and `Large-Disk.gz`.

5.8.3 Special Boot Problems with Kernels from 2.0 Onwards

Problems

Problems booting with LILO might occur after part of the kernel has been loaded. (e. g., a SuSE installation kernel) You can select a kernel at the LILO prompt and this kernel is loaded (some dots are output onto the screen), but starting the kernel fails. Before reaching "uncompressing Linux", the system crashes with different behaviors.

Possible error messages:

- System reboots
- System just hangs
- "crc-error"
- "no free space"
- "Error 0x00"
- "Error 0x01"
- "incomplete literal tree"

Thereafter, access on the floppy is attempted but the system hangs.

Cause

The cause lies in a combination of a big kernel, LILO and faulty hardware. This affects roughly one percent of all machines.

We assume that this is due to a faulty BIOS that has problems with fast memory access.

This problem does *not* occur, if:

- the machine is booted via `loadlin`
- the kernel has been copied to a floppy


```
dd if=/vmlinuz of=/dev/fd0
```

 and booted from there
- a smaller kernel is used, which has been created with


```
make zImage
```

(e. g., an older 1.2.13 kernel) is booted via LILO

The following BIOS settings do not cause any problems either:

- Disable Internal Cache
- **DRAM Precharge Wait State** value: 1 *and*
- **DRAM Wait Burst Timing** value: 0x3333

Solution

First of all, you should be able to install a system. So, if you can neither boot via `loadlin` nor via `setup`, use an old kernel 1.2.13 boot disk for installation. If you do not have such a disk at hand, change the BIOS settings accordingly.

After a successful installation, the question is how to boot in the future. First, use the same medium as during installation. `loadlin` from DOS should not cause any problems. With a boot disk, you should enter the following parameters:

```
load_ramdisk=0 root=/dev/???
```

where ??? is your root partition (e. g., `/dev/hda1`). Thereafter, build your own kernel, since this can be booted using LILO.

5.9 Starting via `loadlin`

Now we want to offer an alternative to boot SuSE Linux, `loadlin`. The program `loadlin` is a DOS program that is capable of booting a Linux kernel from a DOS directory. Thus `loadlin` perfectly integrates itself into an existing DOS/Windows 9x environment. As no entry in the MBR is needed, Windows only notices one or more partitions with unknown ID's. The risk of unwanted side effects due to a Linux installation is thus minimized.

The procedure described below works on both Windows 95 and Windows 98. The files themselves have been written in Windows 95; for this reason we will just talk about Windows 95.

In principle, there are two ways of activating `loadlin`. Switch between various systems via a boot menu or start Linux via `loadlin` from another running system.

Both methods have advantages and disadvantages:

- A boot menu saves the trouble of having to start another operating system first, before you can start Linux.
- You can add other configurations to your boot menu to create a universal starting mechanism.
- You need to modify start files, however, to build a boot menu, which you will have to do by trial and error.
- Changing to Linux from the DOS prompt is very simple.
- A Linux start can be nicely integrated into a Windows 95 session. Double-clicking an icon will start Linux (Windows 95 contains DOS 7.0).



Tip

You should use a boot menu if you are using DOS or Windows 3.x. If using Windows 95, start most easily from the running system. Start menus in Windows 95 are a very complex matter. We are only able to give some hints about them.

5.9.1 Necessary Steps for All loadlin Users

This is what you need to do, whether you decide to use a boot menu or to start from a running system (in DOS, Windows 3.x, or Windows 95):

1. You may have already installed loadlin. If not, do so now using setup.
2. Change to `c:\loadlin` in MS-DOS. There, find the file `linux.par`. Create a file with the editor named `startlin.bat` (you can give it another name) in this directory. Insert the line described in File 5.9.1:

```
c:\loadlin\loadlin @c:\loadlin\linux.par
```

File 5.9.1: Example of a Batch File for Starting Linux

Next, edit the file `linux.par` (File 5.9.2).

```
c:\loadlin\vmlinux # first value must be
                  # the file name of the Linux kernel

root=/dev/xxx    # the device which gets mounted as root FS

ro               # mount root read-only
```

File 5.9.2: Example of the File `linux.par` with Customized Values

Instead of `xxx`, enter your root partition's device name (you wrote down this name in Section 3.4.12 page 58). `startlin.bat` starts Linux. The file `linux.par` is used by `startlin.bat` as well as by `config.sys` and contains essential parameters. Later on, when you are more familiar with Linux, you can add or replace parameters here. If you have built a kernel of your own, just copy it to `c:\loadlin\vmlinux` and, from then on, this kernel will be booted.

5.9.2 Setting up Boot Menus

To configure a boot menu in DOS or Windows 3.x:

1. First, define a boot menu section in your `c:\config.sys` file. Open `c:\config.sys` in an editor and enter something similar to File 5.9.3.

```
[Menu]
menuitem=Win, starting Windows,...
menuitem=DOS, starting DOS, MS-DOS...
menuitem=Linux, starting Linux...
menucolor=15,1
menudefault=Win,5
```

File 5.9.3: Example of First Part of Linux Boot Menu in `c:\config.sys`

Under the label `[Menu]`, define an entry in the boot menu for each OS to boot. Also, define the menu's color and after how many seconds each OS will be automatically started.

2. Below these entries, enter the labels `[Common]`, `[Win]`, `[DOS]`, and `[Linux]`. Commands entered in `[Common]` are always executed. All other entries are OS-specific. See the example in File 5.9.4. As a guide, use the lines in your own `config.sys`. An example may be found in File 5.9.4.

```
[Common]
device=c:\dos\himem.sys /testmem:off
device=c:\dos\emm386.exe noems I=E000-F4FF
dos=high,umb
files=30
buffers=10
shell=c:\dos\command.com

[Win]
devicehigh=c:\dos\dblspace.sys /move
devicehigh=c:\cd\slcd.sys /D:SONY_000 /B:340 /M:P /V /C

[DOS]
devicehigh=c:\dos\dblspace.sys /move
devicehigh=c:\cd\slcd.sys /D:SONY_000 /B:340 /M:P /V /C

[Linux]
shell=c:\loadlin\loadlin.exe @c:\loadlin\linux.par

[Common]
rem Remains blank
```

File 5.9.4: Example of Second Part of Linux Boot Menu in `c:\config.sys`

Now save the file with your changes.

3. Next, edit `c:\autoexec.bat`. Here, you must put the same labels and assign entries to labels, but the notation differs slightly. See the example in File 5.9.5 on the facing page. Notice that the Linux case is not mentioned here because Linux is booted using `loadlin` directly from the `c:\config.sys`

file. The variable %config% contains the selected label (:Win or :DOS). Customize the code to fit your machine.

```
@echo off

rem Entries for all Configurations
switches= /f
set comspec=c:\dos\command.com
prompt $p$g
loadhigh c:\dos\keyb gr,,c:\dos\keyboard.sys
loadhigh c:\dos\doskey
set temp=c:\temp
loadhigh c:\dos\mscdex.exe /D:SONY_000 /E /V /L:H
c:\logimaus\mouse.exe

goto %config%

:Win
c:\dos\smartdrv.exe a- b- c+ 2048 1024
path c:.;d:.;c:\windows;c:\dos;c:\util;
win :
c:\dos\smartdrv /C
goto ende

:DOS
path c:.;d:.;c:\dos;c:\util;
goto ende

:ende
echo * Goodbye *
```

File 5.9.5: Example of autoexec.bat Supporting Linux Boot Menu

4. If you now boot your machine, the boot menu appears and you have five seconds to choose an operating system. Then Windows starts automatically (menudefault=Win,5). If you select 'Linux', Linux starts and awaits your login.

5.9.3 Starting Linux from Within Windows

To create a start icon for Linux to boot Linux from within a running Windows 95 session:

1. Open Windows Explorer. Change to c:\loadlin. With the right mouse button, click on the file startlin.bat and select 'Create Shortcut'.
2. Drag the shortcut onto the desktop.
3. Click on 'Shortcut to startlin.bat' with the right mouse button and select 'Properties'. Go to tab 'Program', click on the button 'Advanced', and click 'MS-DOS mode'. Confirm with 'OK'.

4. Click the 'Change icon' button and select a nice icon. Give the shortcut a suitable name.
5. Double-clicking this new shortcut should bring up a dialog box telling you that Windows 95 is about to switch to DOS mode. If this dialog box bothers you, turn it off in the properties menu.

5.9.4 The Windows Boot Menu

To install a boot menu for Windows 95:

1. You must edit the file `c:\msdos.sys`. First, make the file visible by entering: `C:> attrib -R -S -H c:\msdos.sys`

This is a text file where you have to enter some lines to deactivate the Windows 95 start menu. The [Options] label should resemble File 5.9.6.

```
[Options]
BootGUI=0
BootDelay=0
BootMenu=0
Logo=0
```

File 5.9.6: `msdos.sys` to Start Linux Using a Windows 95 Boot Menu

The parameter `Logo=0` is optional and avoids switching to graphics mode before Windows 95 is started. Booting is much faster and you avoid lots of trouble if you plan to use the DOS emulator in Linux later.

The parameter `BootGUI=0` is for booting Windows 95 directly into DOS mode. To start Windows from this you have to enter `C:> win`, but this is already done by our example `c:\autoexec.bat`, if you have selected Win95 from our menu.

2. Next, define your boot menu in `c:\config.sys`. See File 5.9.7 for an example. Under the Label [Menu], define entries for the boot menu, its color,

```
[Menu]
menuitem=Win95, start Windows 95...
menuitem=DOS, start MS-DOS...
menuitem=Linux, start Linux...
menudefault=Win95,5
```

File 5.9.7: Example `config.sys` (Part One) for Starting Linux Using a Boot Menu with Windows 95

and the delay in starting up automatically.

3. Further down in `config.sys`, enter labels [Win95], [DOS], [Linux], and [Common]. [Common] is for entries that apply every time (this will

very rarely be the case in Windows 95). All other labels are for the corresponding operating system. Use those lines that are already written in your `config.sys`. The example in File 5.9.8 should only be regarded as a hint.

```
[Win95]
dos=high,umb
device=c:\windows\himem.sys /testmem:off

[DOS] device=c:\plugplay\drivers\dos\dwcfgmg.sys
dos=high,umb
device=c:\windows\himem.sys /testmem:off
device=c:\windows\emm386.exe noems I=B000-B7FF
devicehigh=c:\cdrom\torisan.sys /D:TSYCD3 /P:SM

[Linux]
shell=c:\loadlin\loadlin.exe @c:\loadlin\linux.par

[Common]
accdatetime=C+ D+ H+
switches= /F buffers=20
```

File 5.9.8: Example `config.sys` (Part Two) for Starting Linux Using a Boot Menu with Windows 95

Now save this file.

4. Next, edit `c:\autoexec.bat`. Here, the same labels are entered, but the notation is slightly different. The label which has been selected is written to the variable `%config%`. Notice that the Linux case is not included here, because Linux is booted using `loadlin` directly from the `config.sys` file. Your entry should look something like File 5.9.9 on the following page. When you have finished editing the file, do not forget to save it.
5. If you now boot your machine, the Windows 95 boot menu should appear, allowing two seconds to select an entry. If you choose *Command line*, your own boot menu will be displayed. Now you have five seconds to select an operating system. After this delay, Windows 95 starts automatically. If you select 'Linux', Linux starts and awaits your login.

```
@echo off
loadhigh keyb gr,,c:\windows\command\keyboard.sys
goto %config%

:Win95
win
goto ende

:DOS
path c.:d:.;c:\windows\command;c:\util;
loadhigh c:\windows\command\mscdex.exe /D:TSYCD3 /L:x
loadhigh c:\windows\command\doskey
c:\windows\command\mouse.exe
goto ende

:ende
echo * And now? *
```

File 5.9.9: Example autoexec.bat for Starting Linux in Windows 95

6 The X Window System

6.1 Historical Background

The X Window System is the de facto standard GUI for UNIX. Yet the X Window System is far more than this — X11 is a network-based system. Applications running on the machine `earth` can display their results on the machine `sun`, provided the two machines are connected via a network. The network could be a local one (LAN) or a connection between computers thousands of miles away via the Internet.

X11 was first developed as an enterprise of DEC (Digital Equipment Corporation) and the project Athena at MIT (Massachusetts Institute of Technology). The first release of X11R1 was in September 1987. Since release 6, the X Consortium, Inc. has been responsible for the development of the X Window System.

XFree86™ is a freely available implementation of X servers for PC systems. It was developed by a handful of ambitious programmers who founded the XFree86 team in 1992. In 1994, this team went on to found The XFree86 Project, whose aim is to continue research and development on X11 and to provide it to the public. Since March 2000 the completely revised major release XFree86-4.0 has been available for download from <http://www.XFree86.org>. By default, SuSE Linux installs XFree86-4.0. Below we will take a closer look at the features of this version.

SuSE would like to thank the XFree86 team for their help and for their permission to include beta servers on our CD's¹, without which their production would have been much more difficult, if at all possible.

The next sections are about configuring the X server. For this purpose `SaX2`² and `xf86config` will be discussed, simple tools for configuring the X Window System. `SaX2` is intended for configuring XFree86-4.0, `SaX` is used to configure the XFree86 3.3.x versions.


In contrast to the text-based `xf86config`, the two `SaX` versions work directly with the X-server and they can be operated with the mouse. So install, using `YaST`, the program `SaX` (package `sax`, series `x`) or `SaX2` (package `sax2`, series `x`), together with their dependent packages. Whereas XFree86-4.0 contains all the necessary graphics drivers in the base package, if you use XFree86-3.3.x, especially the package `xvga16` — and if it is already known — the X server which matches your graphics card must be installed; the X servers are listed in the series `xsrv`

¹Parts of this documentation are taken from chapter *XFree86 Konfigurieren* from [HHMK96] which was kindly given to us by DIRK HOHNDEL

²`SaX`: *SuSE Advanced X Configuration Tool* The configuration program `SaX2` (`sax2`) to configure XFree86-4.0 makes `XF86Setup` (package `xfsetup`, series `x`) obsolete.

(see Section 3.5.3 page 63). If you have forgotten to select a server, SaX will detect this and install a suitable server, by means of YaST. In XFree86-4.0 the necessary graphics drivers are already included in the base selection.

In order to make optimal use of the hardware available (graphics card, monitor, keyboard), you can optimize the configuration manually. Certain aspects of this optimization will be explained. Even more detailed information on configuring the X Window System can be found in the directory `/usr/share/doc/packages/xfree86` as well as in the manpage for `XF86Config` (man `XF86Config`).



Caution

Be very careful when configuring your X Window System! Never start the X Window System until the configuration is finished. A wrongly configured system can cause irreparable damage to your hardware (this applies especially to fixed-frequency monitors). The authors of this book and SuSE cannot be held responsible for damage. This information has been carefully researched, but this does not guarantee that all methods presented here are correct and will not damage your hardware.

6.2 Version 4.x of XFree86

This version of SuSE Linux comes with version 4.x of XFree86 which differs from the previously used version 3.3 in a number of ways. Overall there are hardly any differences for the user when operating the graphical desktop; applications such as the graphical desktop KDE or GNOME behave with the new version in the same way as version 3.3.6 included in earlier distributions.

What advantages does this version provide?

The new X server is no longer a monolithic program, but just a relatively small basic scaffolding to which the necessary program modules can be later added, if and when required. For example there are no longer many different X servers for different graphics cards, as in the previous version, but just one executable program called XFree86, which can be found in the directory `/usr/X11R6/bin`. This is also the actual X server. The graphics driver, which then takes on the task of controlling the graphics card, is a loadable module.

A similar method is used to support the various input devices, fonts or X protocols. This again consists of individual modules which can be later loaded by the X server. As a rule you don't need to worry about these modules, the configuration of the modules to operate the graphical desktop on your computer is managed as far as possible by SaX2.

Through this module concept, it is easy for a vendor to implement a driver for exotic hardware, such as touchscreens or brand-new graphics cards. The developers have even ensured that the necessary modules for various operating systems

only need to be made available once, which means that a graphics driver module which was compiled in FreeBSD, for example, can also be used in Linux, and vice-versa. This portability, however, is of course limited to the same hardware platform: a module which was compiled for Linux on Power PC's cannot be used on an Intel PC.

Support for the mouse has also been significantly improved. Especially under heavy loads, the reaction of the mouse to mouse movements is considerably faster and more direct than with the previous XFree86 X server. Overall the output speed has also been improved, so that graphics operations are generally performed more quickly than on the old X server due to the completely revised XAA (*XFree86 Acceleration Architecture*).

Compared to XFree86 3.3.x, the configuration file has a slightly different format and is now located in `/etc/X11/XF86Config`: if you want to 'fine tune' your X configuration, details on the structure of the configuration file, and how it functions, can be found in section 6.6 page 154.

Error logging has also been improved. The X server creates a very detailed log file, which you can always find after the X server has started in the file `/var/log/XFree86.0.log`.

One of the further features of this version is the support of special options such as True Type fonts. Other features also include the provision of the 3D protocol extension, **glx**, gamma correction of the screen and the support of multiple graphics cards for **Multihead configurations**. More information on this can be found in Section 6.6 page 154.

What has changed?

XFree86-4.0 is, of course, based on the version 3.3.x. Unfortunately not all drivers are included, as some are very complex and could not be ported to the new XAA architecture. If these graphics cards were supported by SuSE Linux before, you can still use these with the XFree86 version-3.3.x. For such cards, please use the XFree86 version 3.3.6, also contained in SuSE Linux. You can continue to use SaX to configure these cards, as before.

The graphics cards concerned here are those which were until now used together with the following X servers: XF86_S3, XF86_Mach8, XF86_Mach32 and XF86_8514.

In the case of S3 cards, those which previously required the S3 server are not supported by XFree86-4.0, while those which previously worked with the SVGA server will work correctly with XFree86-4.0.

Basically these are the graphics cards S3 Trio3D, Savage4, Savage3D and Savage2000 chips, and almost all S3 Virge cards.

Graphics cards which needed the other X servers listed above (Mach8, Mach32 and 8514) are no longer very common. For these graphics cards — as for all the old S3 cards — the XFree86-3.3 version continues to be available.

6.3 Configuration Using SaX2

The program SaX2 (*SuSE Advanced X Configuration Tool*) provides a simple installation of the X Window System. This is a successor to the tried-and-tested SaX program. SaX was used to configure XFree86-3.3.x, SaX2 is now used for the configuration of XFree86-4.0.

Instructions on how to configure the X Window System with the help of SaX2 can be found in the “Configuration” manual.

6.3.1 Troubleshooting

This section is intended to help you configure the X server if there are problems getting the graphical interface to run at all, or if it runs but with a poor quality display.

Here the main and most frequent problems are discussed which can occur when configuring X with SaX2:

The graphics card is not supported

Using the framebuffer

Unfortunately there are not yet Linux drivers for all graphics cards. If your card is not supported, you can still usually get a graphical display using the *framebuffer* method. If your graphics card is VESA 2 compatible, this method should certainly work.

The advantage of this is that it works with almost every modern graphics card and also practically with every laptop. Because the graphics controller of the card is bypassed, the display will be unaccelerated. The X server simply accesses the *framebuffer* directly after the graphics mode is switched on when the kernel is started.

This works in the following way: when it starts the Linux kernel calls up — even before it switches to 32-bit protected mode — the VGA BIOS of the graphics card and instructs this to switch to a specific VESA graphics mode. The text depiction (that is the text console) then continues in this set graphics mode. The VGA BIOS is written in 16 bit code and therefore cannot be called up while the Linux system is running. As a result, this video mode specified at the start remains until the Linux system is shut down.

In order to now use this VESA framebuffer, support must exist in the kernel for it and the graphics mode must be selected when booting. The SuSE Linux kernel, of course, contains support for the VESA framebuffer. However you need to select the required graphics mode when booting the system. To do this enter the parameter `vga=x` at the LILO boot prompt, where `x` stands for a value to be taken from table 6.1 on the facing page.

You can also specify this parameter directly as a `vga` parameter in the file `/etc/lilo.conf`. The line `vga=x` should be included as a separate line in this configuration file, not as part of the `append` line. The installation of LILO is described in more detail in Section 5.4 page 104.

Desired color depth	Resolution in pixels			
	640x480	800x600	1024x768	1280x1024
256 (8 bit)	769	771	773	775
32768 (15bit)	784	787	790	793
65536 (16bit)	785	788	791	794
16.7 Mill. (24bit)	786	789	792	795

Table 6.1: Possible VESA Modes

After starting the Linux system you should log in to the system as the user `root` and start SaX2 in the following way:

```
earth:/root # sax2 -m 0=fbdev
```

Note that the `0` here is a “zero”, and not a capital “o”.

This instructs the X server to use the driver for the framebuffer. Since resolution, color depth and the frequency rate is fixed, you can — if you are satisfied with mouse and keyboard settings — have SaX2 directly save the automatically detected data and leave SaX2 again.

Using the VGA16 server

If your graphics card is not capable of any of the VESA modes, or if you have an ISA card, you should still be able to configure it to run with the `vga` module. This will be suggested to you if the graphics driver is not automatically detected. Select the `vga` driver in SaX2 in ‘Configure graphics card’ under ‘Driver’. This will limit you in ‘Desktop’, to a resolution of 640x480 pixels with 16 colors (4 bit color depth). If your graphics card supports the VESA framebuffer, the framebuffer method described above is usually preferable, because of this limitation of the `vga` driver. To use the standard `vga` driver straight away when starting SaX2 enter the following command:

```
earth:/root # sax2 -m 0=vga
```

As this parameter only defines which X server SaX2 is to use during configuration, it can be useful to start the X server with this option with a graphics card which is supposed to be supported by XFree86.

In the case of certain “exotic” graphics cards, SaX2 may fail to start in the correct mode initially: if so, do the configuration using the `vga` module, but then choose the necessary options for the required driver module in the ‘Expert’ menu of the ‘Graphics cards’ dialog.

Selecting the mouse directly when starting SaX2

If SaX2 is unable to detect your mouse automatically, you can specify the mouse type when starting SaX2, thus avoiding the automatic mouse detection. For this, there are two command line parameters to determine the mouse protocol and port used. The `-t` parameter is used for defining the protocol used, and the `-n`

parameter for the device file of the mouse. The following example configures an Intellimouse connected to the PS/2 port:

```
root@earth:/root > sax2 -t imps/2 -n /dev/psaux
```

To find out about the available protocols, please refer to the man page of the X server configuration file, **XF86Config**.

Correcting various display errors

In general, SaX2 will configure the graphics card correctly: however there are occasional cases where the display quality is poor. This can be corrected by changing an option in the X server configuration file `/etc/X11/XF86Config`.

This section aims to illustrate errors which occasionally occur when the X server is running and also to provide possible solutions. As a rule the graphic depiction is error-free after configuring with SaX2, but when using certain graphics cards depiction errors may occur which can be removed using an option in the X server configuration file, `/etc/X11/XF86Config`.

Sometimes depiction errors are also caused by hardware limitations; especially if the graphics card uses cheap DRAM memory, it is possible that the bandwidth of the memory is simply no longer adequate for high resolutions with 32 bpp and high frequencies.

Sometimes optimized BIOS settings will get their “revenge”: Linux or the X server does not “like” the BIOS setting **video memory cache mode** which other new BIOS versions provide. Everything functions normally in text mode, but apparently the chip set of the computer then “optimizes” access to the input/output register of the graphics card which can completely mess up the picture with the wrong colors when the X server starts and sometimes cause a total system crash.

If the X server does not start at all, this is almost always because of a syntax or logical error in the configuration file `/etc/X11/XF86Config` and it is possible that the hardware specified in the configuration file cannot be found in the computer. It can be useful to look at the X server logfile and refer also to Section 6.3.2 page 136 when doing this.

Solutions to various possible problems are outlined below.

- Many display problems are caused by the vertical synchronization frequency (`vsync`) being set too high, which results in incorrect values for `hsync` and `DotClock`.

A reduction of the picture frequency rate to 80 Hz, for instance, will still produce a picture which doesn't flutter, and which is also a great deal sharper than one set at 160 Hz. At a picture frequency rate of 80 Hz the pixel rate is only half of that at 160 Hz.

You should, therefore, try reducing the maximum vertical frequency rate used. To do this, choose the ‘Monitor’ tab in SaX and, in the ‘Advanced’ section, set a lower maximum frequency rate. For modern monitors a normal picture frequency rate lies between 80 and 90 Hz.

- Sometimes there are also problems with the “hardware cursor”. You will see a square block or something which looks like a “barcode” where the mouse should be. The solution to this: add the Option `sw_cursor` to the Section `"Device"`.
- The best place to look for information and hints on configuration is in the README files which are sorted by chipset manufacturers’ names in the directory `/usr/X11R6/lib/X11/doc/`. The man pages of the X servers and the XFree86 FAQ are also useful (<http://www.xfree86.org>). The SuSE website also has information at: (<http://www.xfree86.org>) and the SuSE pages for updates:

<http://www.suse.de/en/support/xsuse/>

The following table is ordered in terms of probability, i.e. as a rule you will be successful with step 1, otherwise with the second step. Step 3 is usually a last resort ...

All the options are listed in the device section of the file `/etc/X11/XF86Config`.

- **No mouse cursor, but a colored, square “barcode”.**
 1. Add the option `sw_cursor` to the Section `"Device"`.
 2. Add the option `no_imageblt` or `no_bitblt`.
 3. Add the option `noaccel`.
- **Picture is too narrow or squashed. The monitor settings are already at their limits.**
 1. The frequency rate or the `hsync` frequency is probably set too high and is at the limit of the monitor. Remedy: reduce the `vsync` and/or the `hsync` frequency.
 2. Adjust the picture with the program `xvidtune`. It is possible that just one mode is not quite correct.
 3. Add the parameters `+hsync +vsync` to the modeline and try replacing `+` with `-`.
- **When moving windows, stripes, “lumps” or parts of windows stay where they are. They do not disappear even when the movement is stopped. They only disappear when the desktop is refreshed.**
 1. Reduce the picture frequency rate or the resolution.
 2. Depending on the chip set, use the relevant options from the README files in `/usr/X11R6/lib/X11/doc/`. For example the options `fifo_conservative` or `slow_dram`. *Note:* the options required depend on the particular chipset.
 3. The `noaccel` option, but possibly the options `no_imageblt` or `no_bitblt` may be sufficient.
- **“Noise” — Image interference when moving windows or viewing videos which disappears when the image is stationary.**

1. Reduce the frequency rate, color depth or resolution.
 2. Reduce the frequency rate of the card, or add or remove wait states. This sometimes works with `set_mclk` (not with all chip sets!). More detailed information can be found in the `README` directory. *Attention:* this option is dangerous, the card could be set at too high a frequency).
 3. It is possible that the bus rate is set too highly. Check the bus rate on the PCI/VLB or ISA bus.
- **When starting XFree86, the screen turns black.**
 1. Reduce the frequency rate.
 2. Check the BIOS settings of the computer. Deactivate all “Optimization settings” of the BIOS. Consult your mainboard manual if in doubt. Common culprits are the options **Video memory cache mode**, **AGP Aperture size** and all options which control PCI bus access, such as **PCI Peer concurrency** or similar. You will nearly always find these settings in a menu called **Advanced Chipset Features**.
 3. Consider other possible sources of error: check the system for IRQ conflicts (e. g. PS/2 mouse needing IRQ 12).

6.3.2 The X Server Log File

For analyzing problems with the X server, there is a very detailed log file which the X server creates when it starts. This file is created by the XFree86-4.0 X server according to the following pattern `/var/log/XFree86.Display.Screennummer.log`. If you start just one X server (which is the normal case) and this display is assigned the number “0”, the filename of this log file will usually be called `/var/log/XFree86.0.log`. Note that SaX2 here is an exception to this rule: here, at least temporarily, *two* X servers are running (Display :0) for the configuration dialogs and later a second one (Display :1) to test settings.

You should be aware that the format of this file has changed drastically compared to XFree86-3.3.x. Now a much clearer distinction is made between information messages, values taken from the configuration file, data originating from the computer hardware, and warnings and errors.

SaX2 allows you to view the X server log file with the middle mouse button. Here the various types of messages (error, informal, warning, etc.) are shown in different colors. This helps you to find the problem quickly if the X server does not start.

In general the beginning of such a log file appears as shown in 6.3.1 on the next page.

This provides you with the following information:

This is an XFree86 X server in version 4.0 compatible with X11R6.4 “vendor release 6400”. The release date is 8th March, 2000.

The line **Operating System: Linux 2.2.13 i686 [ELF] SuSE** refers to the system on which the X server was compiled. The kernel version and CPU definition can thus be different from your own system.

```

XFree86 Version 4.0 / X Window System
(protocol Version 11, revision 0, vendor release 6400)
Release Date: 8 March 2000
    If the server is older than 6--12 months, or if your card is
    newer than the above date, look for a newer version before
    reporting problems. (see http://www.XFree86.Org/FAQ)
Operating System: Linux 2.2.13 i686 [ELF] SuSE
Module Loader present
(==) Log file: "/var/log/XFree86.0.log", Time: Sat May 20 13:42:15 2000
(==) Using config file: "/etc/X11/XF86Config"
Markers: (--) probed, (**) from config file, (==) default setting,
        (++) from command line, (!!) notice, (II) informational,
        (WW) warning, (EE) error, (??) unknown.
(==) ServerLayout "Layout[all]"
(**) |-->Screen "Screen[0]" (0)
(**) | | |-->Monitor "Monitor[0]"

```

File 6.3.1: Extract from the X server log file

After these version messages the first login entries appear which the X server creates when it starts. First of all, is this the correct log file? Next to **Time:** the time is specified when the log file was created. Sometimes you might be searching in the wrong log file ; -)

The same thing is valid for the configuration file if you did not specify a different file at the command line, this will always be `/etc/X11/XF86Config` on a normal SuSE Linux system.

The following table 6.2 explains the meaning of the two bracket characters at the beginning of further lines:

Symbol	Meaning
(==)	Defaults of the X server
(--)	Values taken from the system by automatic hardware detection.
(**)	Settings fixed in the configuration file.
(++)	Parameters which you have entered at the command line.
(!!)	Here the X server tells you in detail what it “is doing”.
(II)	Version numbers of X server modules, etc. are usually recorded as “informational messages”.
(WW)	Warnings: here the X server tells you why it is not carrying out certain actions specified in the configuration file or which should be activated by default.
(EE)	Error! These messages lead to the start procedure or the X server crashing. Look out for lines in the log file starting with (EE) if the X server does not start. You can remedy most errors yourself by means of these messages.

Table 6.2: Message types in the X server log file

You can also check the log file within the SaX2 configuration program by pressing the middle mouse button. Errors and warnings of the X server are highlighted

in color. This file is also displayed if the X server which is called up by SaX2 for test purposes does not start or breaks off with an error.

If something unexpected happens when starting SaX2 or during the configuration steps, then all errors and steps concerning SaX2 are logged in the file `/var/log/SaX.log`. X server errors are logged, as described above, in the file `/var/log/XFree86.0.log`. By means of these files you will find clues on how you can carry on from here.

6.3.3 Starting the X Window System

The X Window System is started with the command `startx`. A preconfigured GUI for the `fvwm` window manager is provided for the sample user. We recommend you starting the X Window System from this account, and *not* as `'root'`. X11 server error messages are saved in the `~/X.err` file. The `startx` command has a few options; for instance, you can select 16 bit color depth by typing

```
newbie@earth: > startx -- -bpp 16
```

6.4 Configuration Using SaX

The original version of SaX (*SuSE Advanced X Configuration Tool*) — as opposed to SaX2 — is the configuration tool for the versions 3.3.x of XFree86. If you are using XFree86 version 4.0 you should use SaX2 as described above. If for some reason you need to use XFree86-3.3.x, you should follow the instructions below.



Note

You have to start SaX (`sax`) as user `'root'`. You can also start SaX from YaST: `'System administration'` and then `'Configure XFree86[tm]'` (cf. Section 3.7 page 69).

The program is started from the command line by typing:

```
earth:/root # sax
```

As soon as the program starts it looks for any PCI cards which are installed. If a PCI graphics card is found, it will be identified and displayed in the program under `'graphics card'`.

After the PCI scan, the main window is opened to present the tab windows for the mouse (`'mouse'`), keyboard (`'keyboard'`), graphics card (`'graphics card'`), monitor (`'monitor'`) and desktop (`'desktop'`). (You can change to each window by simply clicking on the appropriate tab). SaX then loads its own hardware database (this may take a few seconds). The data found during the system scan will be presented in its respective categories; the relevant graphics card, for example, can be found in the tab window `'desktop'`.

SaX will do its best to recognize what hardware it can, but to be sure that the configuration settings are correct, you should check them all and modify them if necessary.

Essentially, the program provides you with five “index cards”, i.e. ‘Mouse’, ‘Keyboard’, ‘Graphics Card’, ‘Monitor’ and ‘Desktop’. All cards can be accessed simply by clicking on the appropriate title (“tab” of the card in question).

If your mouse is not yet correctly configured, you can still access SaX via the keyboard. Pressing the **(Tab)** key repeatedly will cycle through each of the entry fields. To change to a different tab window, press **(Tab)** until the desired tab window title is highlighted with a black frame; then you can select the desired tab window using **(←)** or **(→)**, and enter it by pressing **(↵)** (= **(Enter)**). In each tab window there are a number of selection elements, such as buttons, listboxes and entry fields. These can also be accessed by keyboard. To use a button (e.g. ‘Apply’), press **(Tab)** repeatedly until the desired button is highlighted. Pressing **(↵)** or **(Spacebar)** activates the button; the desired action is carried out.

To select an entry in a listbox, push the **(Tab)** button until the desired box is highlighted. By pressing **(↑)** or **(↓)** you can look for an entry by way of its coloring, pressing **(Enter)** will then activate it.

The Mouse

The ‘Mouse’ tab window is the first to appear when the program starts (Figure 6.1).

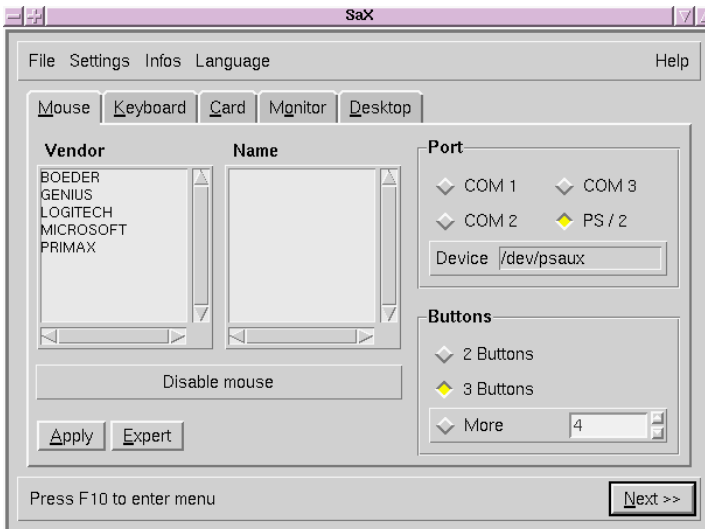


Figure 6.1: SaX: Mouse Settings

If you have already configured your mouse when first installing Linux, e.g. while setting up `gpm`, these settings will be used by SaX and your mouse will work

immediately in the X Window System — then you can ignore this.

If you have not configured your mouse, you should do it now. Press **(Tab)** twice and choose the relevant mouse type in ‘vendor’ (move through the list with **(↑)** and **(↓)**); use **(Enter)** to set the correct vendor name. Go to the selection list using **(Tab)** and choose the relevant type. If you press the ‘Apply’ button you can check to see if your choice was correct. The mouse cursor should move across the screen.

If you’re not sure exactly what mouse type you are currently running on your system, your mouse is not included in the list, or if a serial mouse ‘Microsoft’ cannot run under the ‘Standard Mouse’ protocol, choose the submenu ‘expert’ to set the mouse protocol directly. There you can set further options, such as the baud rate and “three-button emulation”.

Via ‘Expert’ the following tab windows are available:

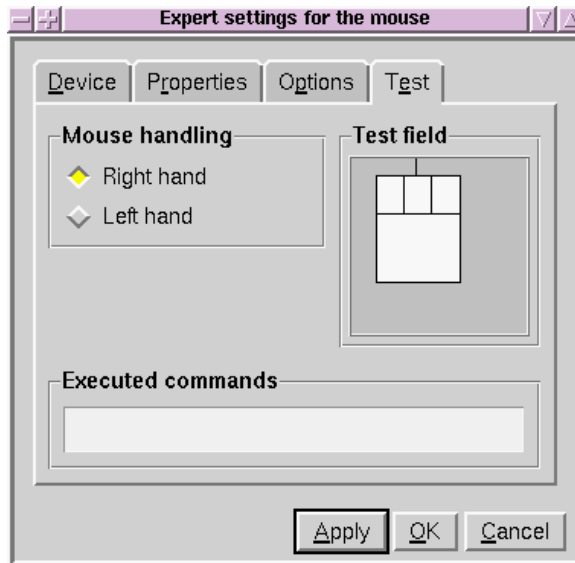


Figure 6.2: SaX: Expert Mouse Settings

‘driver’: If the vendor is not known, then the mouse protocol can be set here. The device file must also be selected. If you have a bus mouse you can try out the corresponding PS/2 variety.

‘Properties’: fine-tuning settings, determining how fast the mouse should react.

‘options’ 3-button emulation, etc.

‘Test’: You can test the mouse configuration using the lower part of the frame ‘Test field’ (Figure 6.2.) If the mouse has been installed correctly, the mouse button symbols on the screen should react when you click on them.

The Keyboard

A Windows 95/98 keyboard with an English keyboard mapping has been set as the default. (Figure 6.3). If you want to use another keyboard, you must enter the correct settings, since the keyboard is one of the few hardware components not to be recognized independently by the hardware scan.

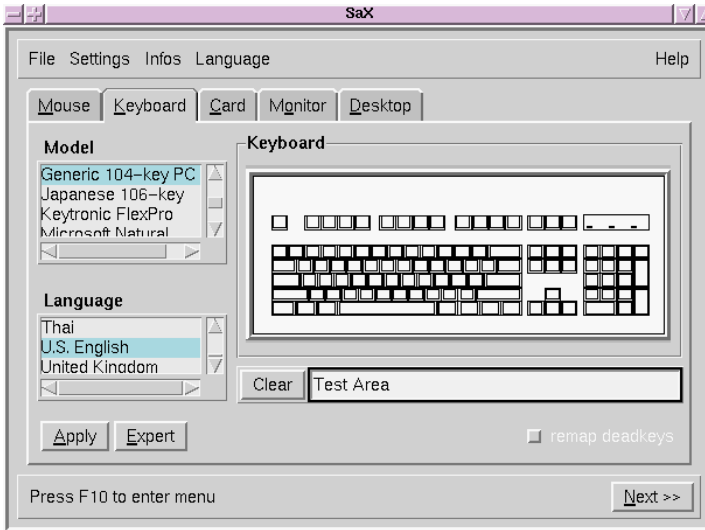


Figure 6.3: SaX: Keyboard

By comparing your keyboard with the ‘Keyboard Image’ shown on the screen you will be able to find the right model connected to your system. Don’t forget to set the ‘language’ to English, if this is not already the case. Next to the delete tab there is a test field where you can check the settings made for the keyboard by entering various characters and seeing if they come out properly.

The switch ‘Use “Dead” keys’ serves with e. g. German keyboards to enable all signs on the keyboard buttons under X.

You probably do not need the settings in ‘Expert’ ...

The changes are put into effect by pressing the ‘Apply’ button.

The Graphics Card

On the ‘card’ tab window, select the card vendors in the left hand list and the card version in the right hand list (Figure 6.4 on the next page). SaX tries to recognize the graphics card independently, which nearly always works for PCI cards. The utility accesses an extensive database of current hardware to achieve this, which can also be found at <http://cdb.suse.de/>. Hardware that has been found will be highlighted in color.

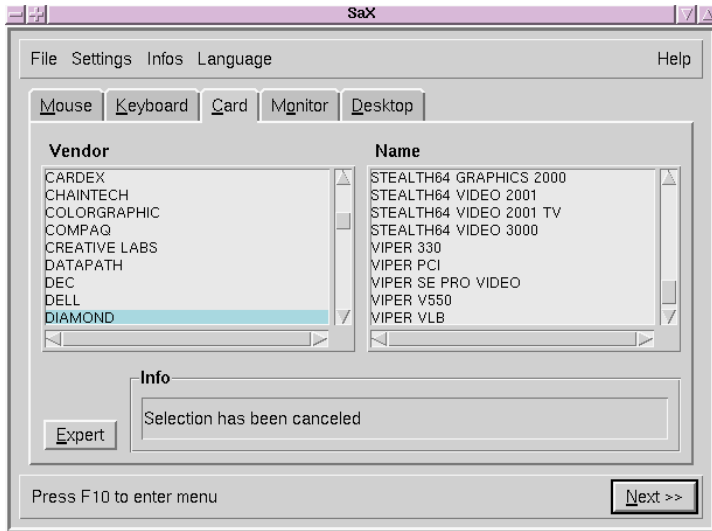


Figure 6.4: SaX: Graphics Card

Note

Sometimes a graphics card of the same series undergoes various “revisions”, in which the hardware has been changed (different Ramdac). So it is possible that SaX may display values which are different to those in the manual of the graphics card. In these cases you should use the details from the hardware vendor’s manual. Select the ‘Expert’ button to do this.

There are advanced options hidden beneath the ‘Expert’ button (Figure 6.5 on the facing page). These are relevant if you choose the X server directly (‘Server settings’). If you want to set memory size or specify the RAMDAC value or if you would like to set a special RAMDAC or clock chip value (in ‘chipsets’), these options are also important. If the screen image looks distorted — if, for example, during a move–window operation, the window fragments or the window title suddenly blink — then you should reduce the RAMDAC value.

Some graphics cards need special ‘options’ which can be found in this expert menu; normally they are not needed.

Tip

ISA cards will not be recognized “automagically”; for these, you have to select the relevant server “by hand”.

If you see the error message “The SVGA Server is not installed...”, you must install the above–mentioned package via YaST (see Section 3.5 page 61).




Figure 6.5: SaX: Graphics Card — Expert Options

The Monitor

The monitor settings are the last great hurdle on your road to a running X server. You find the same divisions in the left tab window list on ‘Monitor’ to choose the monitor ‘Vendor’. Clicking one more option, you can choose your model on the right hand list (‘Type’). Should you be unable to find your own monitor on the list, you can still enter horizontal and vertical frequencies specific to your monitor by pushing the ‘Expert’ button. Normally you should be able to find these in your monitor handbook.

In case you do not have any monitor data available, SaX will set horizontal frequencies to 29–61 kHz and vertical frequencies will be set to 60–70 Hz. Most monitors should not be harmed by these settings.

If the screen remains dark for more than a few seconds after starting the X server or it flickers wildly, you should shut down the server immediately with **(Ctrl) + (Alt) + (←)** (**(←)** denotes the “backspace button”) If you don’t, this might damage to your monitor.

The Desktop

If your graphics card installation was successful, you will have a large number of resolutions and color depths to choose from. They can be administered from the ‘desktop’ menu (Figure 6.6 on the following page).

The ‘desktop’ tab window might remind you of another operating system :-)

In this window you can choose a different resolution (‘resolutions’) for each color depth (‘color’).

If you would like to set a list of several resolutions for a particular color depth, you can do this in expert mode (‘Expert’, Figure 6.7 page 145).

On the tab window ‘Resolution’ you will see:

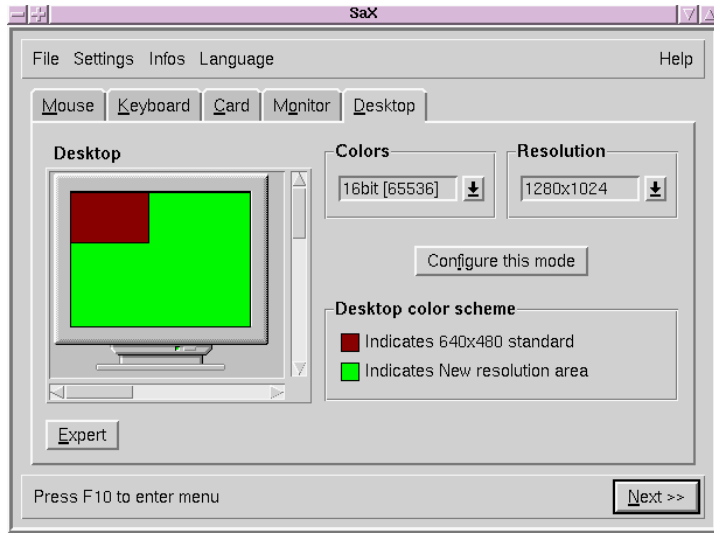


Figure 6.6: SaX: Desktop

‘Resolution’ This tab window is divided into three sections:

‘Possible’ resolutions.

‘Current’ resolutions.

‘Colors’ the list of color depths.

First you should choose your preferred color depth on the vertical button panel to the right (‘Colors’); the resolutions have to be adjusted to these as well.

The X server can start in several color depths, e. g. in 8 bit color depth. 256 colors can be shown on the desktop. Each color depth, in turn, can have different screen resolutions, e. g. 800x600. Not all resolutions are available for each color depth. Thus for some cards a resolution of 1600x1200 pixels at 32-bit color depth is not possible as too little graphics memory is available.

For 8 bit color depth there are several resolutions available, from 1600x1200 down to 640x480; they are listed in the ‘Actual Resolution List’. You can change resolutions by pressing $(\text{Ctrl}) + (\text{Alt}) + (+)$ to cycle forwards through the list or $(\text{Ctrl}) + (\text{Alt}) + (-)$ to go backwards. (you must use the keys $(+)$ and $(-)$ on the number pad). The first entry in the list is always chosen by the X server when it is initialized.

You must copy the desired resolution from the list of ‘possible’ resolutions to the list of ‘current’ resolutions. First you click on the resolution to mark it; then you click on the button ‘I’ to copy the desired resolution to the list of ‘current’ resolutions. By clicking ‘R’ you remove the current resolution.

The sequence of ‘current’ list entries can be changed via the arrow buttons at the lower end of the list window. First you mark the entry you want to

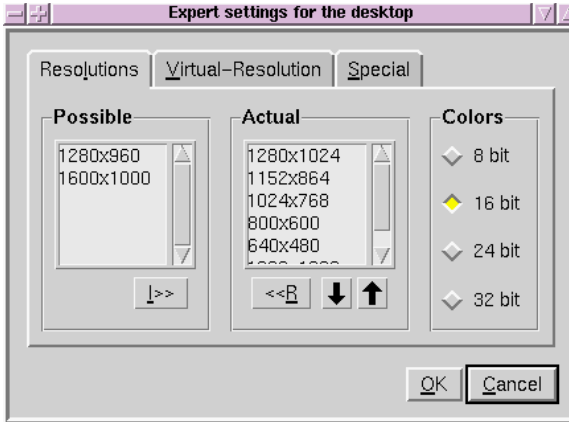


Figure 6.7: SaX: Screen

move and then you left-click the up or down arrow. In this way the entry swaps places with its predecessor or successor.

When you are happy with your settings at 8 bit color depth, you can select another color depth in the box on the right, such as 16 bit. Now you will see a list of ‘possible’ resolutions at this color depth. They can be changed in the same way as described above.

‘Virtual resolution’: this tab window lets you set the virtual resolution of your desktop. If you want to configure a virtual desktop you need to increase the values ‘Virtual X’ (= width) and ‘Virtual Y’ (= height). Generally, this feature is *not* chosen by most people. Background information: the X Window System is capable of defining a virtual desktop that is larger than the actual screen size, for example, using a virtual resolution of 1152x864 with a screen resolution of 800x600.

‘Special’: the ‘special’ tab window allows user-defined entries. In addition you can determine the ‘quality’ of the “modelines”, i.e. you can choose between two different calculation methods.

Now you should set the default color depth with which the X server should start. To do this you need to select the color depth you want in the ‘screen’ tab window (with ‘colors’) as your current color depth; once this is done you should move to the next item on ‘Change current mode’; this should start the X server check ...

Testing the Configuration

After a short time a message box appears; if you are satisfied, click on ‘OK’. Then a background image and a split window should appear; in the left half there is information on your current resolution and on the horizontal and vertical monitor frequencies.

In the right half you should see two button fields, 'size' and 'position', which allow the image to be adjusted. The arrow boxes in 'size' allow the horizontal and vertical image size to be increased or decreased; in the 'position' box you can change the position of the image relative to the monitor. Adjust the screen image to suit you.

Small rectangular image controls can be found in all four corners of the screen image. Ideally they should be all be visible and not display any distortion in color.



Tip

You can only make fairly small monitor adjustments using SaX; it cannot replace manual adjustment using the controls on the monitor itself !

After adjusting the screen image, there are two ways of closing the window:

'save': End your X Window System configuration, saving the current settings. This takes you back to the command line.

'cancel': Interrupt your X server setup and discard the settings.

Press **(Alt) + (F1)** if you want to return to the first console.

6.4.1 Reconfiguring

SaX can also be used to adjust an X server which is already running to your specific needs.

SaX reads, in the existing file `/etc/XF86Config`, where the X Window System stores and analyzes the configuration data. This is why you don't really need to make all the settings from the screen, since SaX adopts the existing, functioning configuration of the X server. SaX displays this data in the tab windows mentioned above.

You are free to do as you want, however, in making new settings for your X server configuration: you can select hardware from a comprehensive monitor database, thereby better adapting the configuration to the monitor's capabilities, as well as adjusting the screen position.

SaX also provides a comfortable GUI to administer color depths and resolutions. It is easily accessed, simply by clicking on the 'screen' tab window.

6.4.2 Troubleshooting

Here are the main and most frequent problems encountered when configuring X windows with SaX:

- If the screen image flickers during the configuration test, or the image turns black, you must shut down the X server immediately, since the monitor could suffer damage if you continue to run the present configuration.

Press **Ctrl** + **Alt** + **←**

You should select the ‘Monitor’ tab window and look for a different monitor, or enter the monitor data by hand; the procedure is the same if the image begins to flicker while the screen is being adjusted.

- In particularly tricky cases, SaX has some command line options available, such as:
 - server vga16**: This will start SaX with the VGA16 server instead of the server appropriate for your card. This server should run on almost all VGA cards. This server is automatically used if your graphics card is not recognized, or if you have an ISA card.

Up-to-date documentation on SaX can be found in the directory `/usr/share/doc/packages/sax`. If something unexpected happens when starting SaX or during the configuration steps, then this will be logged in the files `/root/ServerLog` and `/root/StartLog`. These files will provide you with clues on how best to proceed.

6.5 Configuration Using xf86config

In most cases, XF86Setup is superior to xf86config as a simple configuration tool. There are, nevertheless, some rare cases where XF86Setup fails. If so, just use xf86config, which almost always works.

XFree86-4.0 includes a similar text-based program, xf86config. At some points this contains dialogs which have been somewhat modified and, of course, it writes the configuration file to `/etc/X11/XF86Config`. In XFree86-4.0 however, the use of xf86config is usually not needed, since “problem” graphics cards can also be configured with the “Framebuffer” or with the **vga** module as described in Sections 6.3.1 page 132 and 6.3.1 page 133.

Make sure you have the following information available:

- mouse type, port to which the mouse is connected and baud rate (the baud rate is normally optional).
- specifications of the graphics card
- monitor data (frequencies, etc.)

If these settings are known, or you have your manuals at hand, you can start configuring. Remember that you have to be ‘root’ to do this.

The configuration is started with:

```
earth:/root # /usr/X11R6/bin/xf86config
```

Mouse

After the welcome screen, you are asked about your mouse type. You are offered the following selections (see Output 6.5.1 on the following page:

1. Microsoft compatible (2--button protocol)
2. Mouse Systems (3--button protocol)
3. Bus Mouse
4. PS/2 Mouse
5. Logitech Mouse (serial, old type, Logitech protocol)
6. Logitech MouseMan (Microsoft compatible)
7. MM Series
8. MM HitTablet

Output 6.5.1: Mouse Selection for X

While selecting the mouse, you should bear in mind that many of the new Logitech mice are either Microsoft compatible, or use the MouseMan protocol. The selection **Bus Mouse** refers to any bus mouse, including Logitech!

Selection is made by entering the relevant number. There may be a question whether “ChordMiddle” should be activated. This is necessary for some Logitech mice or trackballs, to activate the middle mouse button.

Please answer the following question with either ‘y’ or ‘n’.

Do you want to enable ChordMiddle?

If you have a two-button mouse, you can emulate the third button by answering ‘y’ to the next question.

Please answer the following question with either ‘y’ or ‘n’.

Do you want to enable Emulate3Buttons?

The middle button is emulated by simultaneously pressing the two mouse buttons.

Next you must specify the mouse’s interface:

Now give the full device name that the mouse is connected to, for example /dev/tty00. Just pressing enter will use the default, /dev/mouse.

Mouse device:

If you have already entered a port for your mouse during the system installation, just enter **/dev/mouse**.

Keyboard

Next you are asked whether to assign **Meta** (ESC) to the left **(Alt)** key and to assign **ModeShift** to the right **(Alt)** key.

Please answer the following question with either ‘y’ or ‘n’.


Do you want to enable these bindings for the Alt keys?

If you answer ‘y’, you can access the German keyboard via the right **(Alt)**, and the left **(Alt)** can serve as the meta key.³

³e. g., in Emacs.

Monitor

Next you should specify your monitor. You should be extremely careful with vertical and horizontal frequencies! These can be found in your monitor handbook.



Caution

Setting frequencies incorrectly can lead to irreparable damage to your monitor! The X Window System only addresses video modes which operate the monitor in the given frequency range. Entering frequencies for which the monitor was not designed can cause severe damage to it!

Some monitors are listed under `/usr/X11R6/lib/X11/doc/Monitors`⁴.

To enter the horizontal frequency, the following selection is displayed (see Output 6.5.2):

```
hsync in kHz; monitor type with characteristic modes
 1 31.5;                Standard VGA, 640x480 @ 60 Hz
 2 31.5 - 35.1;        Super VGA, 800x600 @ 56 Hz
 3 31.5, 35.5;         8514 Compatible, 1024x768 @ 87 Hz interl.
                       (no 800x600)
 4 31.5, 35.15, 35.5;  Super VGA, 1024x768 @ 87 Hz il.,
                       800x600 @ 56 Hz
 5 31.5 - 37.9;        Extended Super VGA, 800x600 @ 60 Hz,
                       640x480 @ 72 Hz
 6 31.5 - 48.5;        Non-Interlaced SVGA, 1024x768 @ 60 Hz,
                       800x600 @ 72 Hz
 7 31.5 - 57.0;        High Frequency SVGA, 1024x768 @ 70 Hz
 8 31.5 - 64.3;        Monitor that can do 1280x1024 @ 60 Hz
 9 31.5 - 79.0;        Monitor that can do 1280x1024 @ 74 Hz
10 Enter your own horizontal sync range
Enter your choice (1--10):
```

Output 6.5.2: Entry for the Monitor's Horizontal Frequency

You should only choose one of the predefined modes if you are unsure of the settings for your monitor. Selection '10' allows you to enter your own frequencies. The next screen asks you to enter your monitor's vertical frequency (see Output 6.5.3 on the following page). Again, using the known values (i.e., choice '5') is preferable to using one of the items '1' to '4'.

Next you should enter a name, vendor name and model for your monitor:

```
Enter an identifier for your monitor definition:
Enter the vendor name of your monitor:
Enter the model name of your monitor:
```

These are just descriptive names used to document your configuration which do not affect the configuration itself. Merely pressing **Enter** will select the default values which are usually sufficient.

Your monitor configuration is now complete.

⁴We cannot be held liable, of course, if this information is inaccurate!

```
1 50-70
2 50-90
3 50-100
4 40-150
5 Enter your own vertical sync range
```

Enter your choice (1--5):

Output 6.5.3: Vertical Frequency Choices

Graphics Cards / X server

Next you must specify your graphics card:

Do you want to look at the card database?

If you enter 'y', a selection of predefined cards is presented.

Here you can select your card by pressing the corresponding number. Do not trust this list blindly, since there can be differences in clock chip and RAMDAC⁵ settings!

This is why there is a menu item later to select a RAMDAC and a clock chip, even though you have entered them already. Then the predefined settings for this card will be presented as an extra option.

The card definitions contain information on clock chips, RAMDAC and the X server to be used. Furthermore, some valuable information concerning the card is written to the device section in `XF86Config`.

If your card is not listed, do not panic. You should switch back to the normal configuration by selecting 'q'. Only select one of the defined cards if it matches your card exactly! Selecting a card with a similar name is not recommended. Similar names do not necessarily refer to similar hardware.

Further information on how to configure your card is given in chapter Section 6.6 page 154.

After specifying your card, the X server is next. `xf86config` displays the choices, as seen in Output 6.5.4 on the facing page.

Choice '5' only appears if you have selected one of the predefined cards in the previous step. In this case, choose '5' to select the X server most suitable for your card.

When you have selected a server, you are asked if you want to create a symbolic link to `/usr/X11R6/bin/X`. If you answer with 'y', you are asked whether you want to put it in `/var/X11R6/bin/X`.

Do you want to set it in `/var/X11R6/bin`?

Reply with 'y', since it may not always be possible to write to `/usr`.

Afterwards, if you have selected '4' (the accelerated servers) in the previous selection, a menu is presented of all available accelerated X servers, as shown in Output 6.5.5 on the next page.

After selecting your X server, you now have to configure your graphics. First you should specify the amount of memory the card has, as seen in Output 6.5.6 on the facing page.

⁵Random Access Memory Digital-to-Analogue Converter.

- 1 The XF86_Mono server. This a monochrome server that should work on any VGA--compatible card, in 640x480 (more on some SVGA chipsets).
- 2 The XF86_VGA16 server. This is a 16-color VGA server that should work on any VGA--compatible card.
- 3 The XF86_SVGA server. This is a 256 color SVGA server that supports a number of SVGA chipsets. It is accelerated on some Cirrus and WD chipsets; it supports 16/32-bit color on certain Cirrus configurations.
- 4 The accelerated servers. These include XF86_S3, XF86_Mach32, XF86_Mach8, XF86_8514, XF86_P9000, XF86_AGX, XF86_W32 and XF86_Mach64.

These four server types correspond to the four different "Screen" sections in XF86Config (vga2, vga16, svga, accel).

5 Choose the server from the card definition, XF86_S3.

Which one of these four screen types do you intend to run by default (1--4)?

Output 6.5.4: Selecting an X Server

Select an accel server:

- 1 XF86_S3
- 2 XF86_Mach32
- 3 XF86_Mach8
- 4 XF86_8514
- 5 XF86_P9000
- 6 XF86_AGX
- 7 XF86_W32
- 8 XF86_MACH64

Which accel server:

.

Output 6.5.5: Accelerated X Server Options

How much memory do you have on your graphics card:

- 1 256K
- 2 512K
- 3 1024K
- 4 2048K
- 5 4096K
- 6 Other

Enter your choice:

Output 6.5.6: Selecting Video Memory

Next, you must enter the name, vendor name and type for your graphics card. These are merely descriptive entries. If you earlier selected a card from the predefined list, pressing **(Enter)** will enter this as the default.

Enter an identifier for your graphics card definition:

Enter the vendor name of your graphics card:

Enter the model (board) name of your graphics card :

If you chose an accelerated X server, you must enter the RAMDAC settings. This only applies to the S3 and AGX servers.

In most cases, simply pressing **(Enter)** will suffice. If you have selected a graphics card that supports a specific RAMDAC, this should be chosen here (see Output 6.5.7).

```

1 AT&T 20C490 (S3 server)          att20c490
2 AT&T 20C498/21C498/22C498 (S3)  att20c498
3 AT&T 20C505 (S3)                att20c505
4 BrookTree BT481 (AGX)           bt481
5 BrookTree BT482 (AGX)           bt482
6 BrookTree BT485/9485 (S3)       bt485
7 Sierra SC15025 (S3, AGX)        sc15025
8 S3 GenDAC (86C708) (autodetected) s3gendac
9 S3 SDAC (86C716) (autodetected)  sdac
10 STG-1700 (S3)                  stg1700
11 TI 3020 (S3)                   ti3020
12 TI 3025 (S3)                   ti3025

```

Output 6.5.7: Setting a RAMDAC

After answering this question, you can enter a clock chip for accelerated cards, if you have one (see Output 6.5.8). Entering a clock chip avoids clock lines, as the clocks needed can be programmed.

```

1 AT&T 20C490 (S3 server)          att20c490
2 AT&T 20C498/21C498/22C498 (S3)  att20c498
3 AT&T 20C505 (S3)                att20c505
4 BrookTree BT481 (AGX)           bt481
5 BrookTree BT482 (AGX)           bt482
6 BrookTree BT485/9485 (S3)       bt485
7 Sierra SC15025 (S3, AGX)        sc15025
8 S3 GenDAC (86C708) (autodetected) s3gendac
9 S3 SDAC (86C716) (autodetected)  s3_sdac
10 STG-1700 (S3)                  stg1700
11 TI 3020 (S3)                   ti3020
12 TI 3025 (S3)                   ti3025
13 Normal DAC                      normal

```

Output 6.5.8: Setting the Clock Chip

If a card without a clock chip is selected, just press **(Enter)** (thus not selecting a clock chip). If a card has been selected, the clock chip is set as default (if there is one).

If no clock chip has been set, `xf86config` suggests running `X -probeonly` to determine the clock timings supported. These are automatically written in `XF86Config` in a separate `clocks` line.

Here, we must explain why the automatically defined settings can be *really dangerous*: if the card has a programmable clock chip, the X server, when probing, cannot distinguish between the different clocks and only recognizes clocks 0, 1, and sometimes, 2. All other values are more or less random numbers (normally, clocks 0, 1, and 2 are repeated and are replaced by zeros).

All clocks apart from 0 and 1 are strongly influenced by the preprogrammed clock chip. Thus, clock 2 could have a different setting when probed (and which was written to the file `XF86Config`) than when the X server is later started. Then all the timings would be wrong and the monitor could be severely damaged!

A good indication of a programmable clock chip (and the problems this might entail) are many zeros or repeated timing values. Never ever write such values to `XF86Config`!

To configure clock chips, follow these steps:

- The best way is to enter an existing (*programmable*) clock chip if there is one. It will be programmed accordingly and your `XF86Config` will not contain clock lines. You can compare chips on the card with the chips offered in the menu. Most newer S3 cards have a programmable clock chip.
- If you *do not have a programmable* clock chip, you should launch X -probeonly and compare these values with those in the manual. If these values correspond (± 2), you should enter them in `XF86Config`. If there are no hints in the manual, you can determine the values by running X -probeonly (this works best on an unloaded machine). Check whether the values are correct, since clock values cannot be determined for every card. (Many zeros or repeating values are a sign of invalid settings.) Enter the correct values into `XF86Config`. Do not omit any values; do not try to rearrange them or change them in any way. The values have to be entered in their exact order.


Exception: if the P9000 server is used, the order is irrelevant; just enter the modes for the desired clock in the *clocks line*.

- In general: if there is a programmable clock chip, there should be *no* clocks line in `XF86Config` (exception, P9000).
For cards without a programmable clock chip, there should be a *clocks line* in `XF86Config`. This avoids the tedious (and sometimes even dangerous) testing at each startup. Furthermore, for cards with unreadable values, there are no invalid values and there is no risk to your monitor.

After having read the previous section, if you want to let clocks be recognized automatically, just answer 'y' to the following question:

```
Do you want me to run 'X -probeonly' now?
```

Now the screen will turn black and then a list of probed clocks will be presented, or a message will appear that no clocks could be found. If you have selected a clock chip, this question will not appear, since the clocks are then programmed automatically. In this case, this section is skipped.



Caution

If the previous question has been answered with 'y' and the screen remains black for more than 30 seconds, you should cancel testing immediately with **(Ctrl) + (Alt) + (←)**, or **(Ctrl) + (c)**. If this does not work, switch off the monitor and the computer to prevent the hardware from being damaged!

Saving your Configuration

Now the configuration file has to be written. It is recommended that you write it to `/etc/XF86Config` to ensure that, even in a networking environment, each machine has its own configuration file — even if they share the `/usr` file system. `xf86config` first suggests that it should write `XF86Config` to the current directory. You should answer 'no' to this:

```
Do you want it written to the current directory as 'XF86Config'?
```

Then you are asked where to save your configuration:

```
Please give a path + filename to write to:
```

Here, you should enter: `"/etc/XF86Config"`.

Now `xf86config` exits to the command line. This completes the configuration of the X Window System.

6.6 Optimizing the Installation of the X Window System

This section describes the configuration file, `/etc/X11/XF86Config`. Each *section* starts with the keyword **Section** `<name of section>` and ends with **EndSection**. Below there is a rough outline of the most important sections.

Afterwards you will learn how to integrate additional fonts, how you can configure input devices and how 3D acceleration is implemented. This is also managed in certain sections of the `XF86Config` file, of course, although integrating an additional font requires the help of external programs, which are included with SuSE Linux or are part of the default installation. The methods discussed here aim to illustrate the possibilities available and serve as an incentive, but they do not claim to cover all eventualities.

The programs `SaX2` and `xf86config` (for XFree86-4.0) create the file `XF86Config`, by default in `/etc/X11`. This is the primary configuration file for the X Window System. You can find all the settings here concerning your graphics card, mouse and monitor.

`XF86Config` is divided into several sections, each one dealing with a certain aspect of the configuration. A section always has the same form:

```
Section <name of section>
  entry 1
  entry 2
  entry n
EndSection
```

The following types of sections exist:

Files	This section describes the paths used for fonts and the RGB color table.
ServerFlags	General switches are set here.
InputDevice	Input devices are configured in this section. In contrast to XFree86-3.3, both keyboards and mice as well as special input devices (touch pad, joysticks etc.) are configured via this section. Important terms here are Driver and the options which are defined by Protocol and Device .
Monitor	Describes the monitor used. The individual elements of this are the name, which is referred to later in the Screen definition, the bandwidth and the allowed sync frequencies (HorizSync and VertRefresh). Settings are given in MHz, kHz and Hz. Normally, the server refuses any modeline that does not correspond with the specification of the monitor. This is to prevent too high frequencies from being sent to the monitor by accident.
Modes	The modeline parameters are stored here for the specific screen resolutions. These parameters can be calculated by SaX2 on the basis of the values given by the user and normally do not need to be changed. You can intervene manually at this point, however, if, for example you want to connect a fixed frequency monitor. An exact explanation of the individual parameters would be too much for this book. You can find details on the meaning of individual number values, however, in the HOWTO file <code>/usr/share/doc/howto/en/XFree86-Video-Timings-HOWTO.gz</code> .
Device	This section defines a specific graphics card. It is referenced by its descriptive name.
Screen	This section puts together a Driver (e.g., vga2), a monitor and a Device to form all the necessary settings for XFree86. In the Display subsection you can specify the size of the virtual screen (Virtual , the ViewPort and the Modes) used with this virtual screen.
ServerLayout	This section defines the layout of a single or multihead configuration. The input devices InputDevice and the display devices Screen are combined into one section.

Table 6.3: Sections in `/etc/X11/XF86Config`

We will now take a closer look at **Monitor**, **Device** and **Screen**. Further

information on the other sections can be found in the manpage for **XFree86** (man **XFree86**) and manpage for **XF86Config** (man **XF86Config**).

There can be several different **Monitor** sections in **XF86Config**. Even multiple **Screen** sections are possible; which one is started depends on the server started.

Screen Section

First we will take a closer look at the screen section. As mentioned above, this combines a monitor with a device section and determines which resolution using which color depth, should be used.

A screen section might look like the example in File contents 6.6.1.

```
Section "Screen"
    Driver      "accel"
    Device      "Miro Crystal 40SV"
    Monitor     "EIZO T563-T"
    DefaultColorDepth 16
    Subsection "Display"
        Depth      8
        Modes      "1024x768" "800x600" "640x480"
        ViewPort   0 0
        Virtual    1024 768
    EndSubsection
    Subsection "Display"
        Depth      16
        Modes      "1280x960" "1152x864" "1024x768" "800x600"
        ViewPort   0 0
        Virtual    1280 960
    EndSubsection
    Subsection "Display"
        Depth      32
        Modes      "1024x768" "800x600" "640x480"
        ViewPort   0 0
        Virtual    1024 768
    EndSubsection
EndSection
```

File 6.6.1: An example of the screen section in `/etc/XF86Config`

The line **Identifier** (here **Screen[0]**) gives this section a defined name with which it can be uniquely referenced in the following **ServerLayout** section.

The lines **Device** and **Monitor**, specify the graphics card and the monitor which belong to this definition. These are just links to the **Device** and **Monitor** sections with their corresponding names or “identifiers”. These sections are discussed later in more detail.

Using **DefaultColorDepth**, you can select which color depth mode the server will start with if this is not explicitly stated.

There is a **Display** subsection for each color depth. **Depth** assigns the color depth which is valid for this subsection. Possible values for **Depth** are: 8, 16, 24 and 32. Not every X server supports all these modes. For most cards, 24 and

32 are basically the same, some take 24 for packed pixel 24bpp mode, whereas others choose 32 for padded pixel mode.

After the color depth, a list of resolutions is set (**Modes**). This list is checked by the server from left to right. For each resolution, a suitable **Modeline** is searched for, which has to correspond to one of the given clock rates or a clock rate to program the card.

The first resolution found is the so-called **Default mode**. With **(Ctrl) + (Alt) + (gray +)**, you can switch to the next resolution in the list to the right, with **(Ctrl) + (Alt) + (gray -)**, to the left, thus enabling you to vary the resolution while X Windows is running.

The last line of the **Display** subsection with **Depth 16** refers to the size of the virtual screen. The maximum possible size of a virtual screen depends on the amount of memory installed on the graphics card and the desired color depth, not on the maximum resolution of the monitor. Since modern graphics cards have a large amount of video memory you can create very large virtual desktops. You should note, however, that you may no longer be able to use 3D functionality if you fill practically the entire video memory with a virtual desktop. If the card has 16 MB video RAM, for example, the virtual screen can be up to 4096x4096 (!) pixels in size at 8-bit color depth. Especially for accelerated cards, however, it is not recommended to use up all your memory for the virtual screen since this memory on the card is also used for several font and graphics caches.

Device Section

A device section describes a specific graphics card. You can have as many device entries in `XF86Config` as you like, as long as their names are differentiated, using the keyword **Identifier**. As a rule — if you have more than one graphics card installed — the sections are simply numbered in order the first one is called **Device[0]**, the second one **Device[1]** etc. In the file 6.6.2 you can see the section from the **Device** section of a computer in which a Matrox Millennium PCI graphics card is installed.

```
Section "Device"
    BoardName      "MGA2064W"
    BusID          "0:19:0"
    Driver         "mga"
    Identifier     "Device[0]"
    VendorName    "Matrox"
    Option        "sw_cursor"
EndSection
```

File 6.6.2: The device section of the file `/etc/X11/XF86Config`

If you use **SaX2** for configuring, then the device section should look something like the above diagram. Both the **Driver** and **BusID** of course are dependent on the hardware installed in your computer and are detected by **SaX2** automatically. The **BusID** defines the PCI or AGP slot in which the graphics card is installed. This matches the ID which is displayed by the command `lspci`. Note here that the

X server wants details in decimal form, while `lspci` displays these in hexadecimal form!

Via the **Driver** parameter you can specify the driver to be used for this graphics card. If the card is a Matrox Millenium then the driver module is called **mga**. The X server then searches through the **ModulePath** defined in the **Files** section in the `drivers` subdirectory. In a standard installation this is the directory `/usr/X11R6/lib/modules/drivers`. For this purpose simply `_drv.o` is added to the name, so in the case of the **mga** driver the driver file `mga_drv.o` is loaded.

The behavior of the X server or of the driver can also be influenced through additional options. An example of this is the option `sw_cursor` which is set in the device section. This deactivates the hardware mouse cursor and depicts the mouse cursor using software. Depending on the driver module, you have various options available which can be found in the description files on the driver modules in the directory `/usr/X11R6/lib/X11/doc`. Generally valid options can also be found in the manpage for **XF86Config** (`man XF86Config`) and manpage for **XFree86** (`man XFree86`).

Monitor Section

Monitor sections each describe, in the same way as the device sections, one monitor. The configuration file `/etc/X11/XF86Config` can again contain as many **Monitor** sections as you want in `XF86Config`. The server layout section specifies which monitor section is relevant.

The same applies for monitor definitions: they should only be set by experienced users. A critical part of the monitor section is the modelines which set horizontal and vertical timings for the appropriate resolution. The monitor properties, especially the allowed frequencies, are stored in the monitor section.

Caution

Unless you have an in-depth knowledge of monitor and graphics card functions, nothing should be changed in the modelines, since this could cause severe damage to your monitor!



For those who may want to develop their own monitor descriptions, the documentation in `/usr/X11/lib/X11/doc` might come in handy. The section **[FCR93]** deserves a special mention. It describes in detail how the hardware functions and how modelines are created.

Luckily a “manual” setting of the modelines is hardly ever needed nowadays. If you are using a modern multisync monitor, the allowed frequencies and optimal resolutions can, as a rule, be read directly from the monitor by the X server via DDC, as described in the `SaX2` configuration section. If this is not possible for some reason, you can also use one of the VESA modes included in the X server. This will function with practically all graphics card and monitor combinations.

6.6.1 Inserting Additional (True Type) Fonts

A standard X11R6 X server installation also includes a large number of fonts. These can be found in the directory `/usr/X11R6/lib/X11/fonts`, each divided into logically connected groups in subdirectories. Make sure that only subdirectories of the X server are used which —

- are entered in the files section, **Files** of the file `/etc/X11/XF86Config` as **FontPath**.
- contain a valid `fonts.dir` file.
- were not closed while the X server was running using the command `xset -fp`.
- or which were started while the X server was running using the command `xset +fp`.

Since version 4.0, XFree86 can use not only its own format **Type1** (a Postscript format) for scalable fonts and **pcf** for bitmap ones, but also the **ttf** (True Type font) of this file format. As described in Section 6.2 page 130, this support is provided via loadable modules of the X server. Thus you can also use directories containing True Type fonts together with the X server. To do this, hardly any preparation is needed.

A big advantage of most True Type fonts, apart from their very good scalability, is that these fonts almost always contain more than the normal 255 characters of the font for western Europe coded in “iso-8859-1”. With these fonts you can display Cyrillic, Greek or eastern European languages without any problem, and with special software, even Asian languages.

This description is essentially about the use of fonts as 8-bit character sets. If you want to use characters of Asian languages (Japanese, Chinese, etc.), you can use special editors which are also available in SuSE Linux.

An 8-bit character set contains 255 characters and basically consists of the US-ASCII character set, which defines only the first 128 of 255 possible characters, to be able to expand further characters. One text character thus occupies 8-bits in the computer memory. As 127 characters are certainly not enough to record the special characters, for example, of all European languages, the various languages are combined into groups and this group is then given a short name. The relevant character set is named according to the appropriate norm as the “iso-8859-x” character set, where by the x stands for a number from 1 to 15. The exact order of characters in the `iso-8859-1` character set can be found from the manpage for `iso-8859-1` (`man iso-8859-1`).

The more well-known codings are listed in table 6.4 on the following page : further ones can be taken from the above-mentioned manual page.

The user must then — depending on the language used — select the matching coding. Especially when transferring texts between different computers, the coding used must also be transferred. The advantage of this procedure is obvious: To receive support for regional special characters you only need to select the correct coding and immediately (nearly) all programs will be able to portray these special characters, since (nearly) all programs use an 8-bit value (one byte) to

Font	Supported regions, contains special characters
iso-8859-1	West European languages: Spanish, German, French, Swedish, Finnish, Danish and others
iso-8859-2	Central and Eastern Europe: Czech, Rumanian, Polish, German and others
iso-8859-5	Cyrillic characters for Russian
iso-8859-7	Greek characters for Greek
iso-8859-15	As iso-8859-1, but with characters for Turkish and the Euro sign.

Table 6.4: Important Font Codings

represent a text character. If the wrong coding is chosen then the special characters will be wrongly depicted. With most X applications, as well as with the KDE desktop, you can usually select the coding of the character set when you are configuring the font to be used. In X applications the coding is usually referred to as **Encoding**.

The disadvantage of this method is that some language combinations are impossible: You can not, for example, easily write a German text with umlauts in which you mention Russian place names in Cyrillic.

This dilemma can only be solved using a different approach, with the use of Unicode. Unicode codes characters — differently to ASCII — not with one but with two or even more bytes, allowing considerably more characters to be represented. Only if you use Unicode can you depict Asian languages with more than 127 characters, such as Chinese, Japanese or Korean, on the computer. The disadvantage of this method is that most existing software cannot handle these characters and that you can only read or write texts yourself with Unicode characters using special software. If you want more information on using Unicode fonts in Linux, take a look at the URL <http://www.unicode.org>. It is expected that in future more and more programs will support Unicode characters. In SuSE Linux there is the program `yudit` to enter texts in Unicode. The program `yudit` can be found in package `yudit`, series `xap` or after installation via the SuSE menu, under **Office** and then under **Editors**.

After these observations, we now have a step by step description of the installation of additional fonts, using the example here of True Type fonts.

Locate the fonts that you want to install in your X Window System. If you have licensed True Type fonts then you can simply use these on your system. Mount the partition containing these fonts.

You should create a font directory — if this does not yet exist — and change to it. SuSE Linux already has a directory called `/usr/X11R6/lib/X11/fonts/truetype`, you can copy the relevant fonts to this directory.

```
earth:/root # cd /usr/X11R6/lib/X11/fonts/truetype
```

Create links to the ttf files and create the font directory. Note that for True Type fonts you need a special program called `ttmkfdir`, package `ttmkfdir`, series `xap` to create the file `fonts.dir`. Traditional X fonts are included using the com-

mand `mkfontdir`. Instead of the path `/path/to/the/fonts`, set the corresponding path in which these fonts are located.

```
earth:/usr/X11R6/lib/X11/fonts/truetype #
ln -s /pfad/zu/den/fonts/*.ttf .
earth:/usr/X11R6/lib/X11/fonts/truetype #
ttmkfdir -o fonts.dir
```

If the X server is already running, you can now make the fonts dynamically available. To do this enter:

```
newbie@earth:/home/newbie > xset +fp /usr/X11R6/lib/X11/fonts/truetype
```



Tip

The `xset` command accesses the X server via the X protocol. It must therefore have access permissions for the X server currently running. This is the case, for example, if the user `newbie` has started the X server. You can find more on this in the manpage for `xauth` (`man xauth`)

To set up the fonts permanently you should add this search path to the file `XF86Config`. You can use `SaX2` to do this. To change the fonts path you must select the ‘Custom’ configuration mode of `SaX2`. In ‘Path dialog’ you can add the directory, with ‘Add’, to the directories already listed.

You should test if the fonts were set up correctly. To do this use the command `xlsfonts`. If the fonts are correctly installed, the list of all installed fonts, including the newly installed True Type Fonts, is displayed. You can also use the KDE program, `kfontmanager`, which displays the installed fonts with an example text.

```
newbie@earth:/home/newbie > xlsfonts
```

These newly installed fonts can then be used in all X applications.

6.6.2 Setting Up Input Devices

If you wanted to describe all the possible settings of input devices you could write a whole book on it, simply because there are a huge number of input devices available in the computer world.

The following section will try to give you a few recipes with which you can configure these input devices. The somewhat abstract term “input devices” refers to the mouse and keyboard, but also to touch pads, etc. Look for the relevant heading below to find more information on your relevant topic.

You can, of course, find more detailed information in the X server documentation, especially in the manpage for `XF86Config` (`man XF86Config`).

Mouse

Wheel mice If you have set up the mouse with `SaX2`, then you can also use the basic functionality. You might have a mouse with a wheel (a so-called “wheel mouse”), in which case you would like to use this wheel in X applications, e.g.

for scrolling. If you have set up such a wheel mouse with SaX2 using the extended settings in the SaX2 mouse menu, 'wheel mouse', then movements on the wheel are passed on to applications as movements on buttons 4 and 5. Unfortunately, only a few X applications can handle this. To be able to use the wheel in all programs, you just need to install the program package `imwheel`, series `ap`, in the running X server.

You can do this in the file `~/.xinitrc`. You can insert the program command, for example directly under the commentary line `finally start the window manager`. This small utility converts the "mouse clicks" which are produced by moving the wheel, to configurable key clicks. These are preconfigured to `(PgUp)` and `(PgDn)`, the program can also be configured via the file `/etc/imwheelrc`. In this way you can operate every X program which can be used via the keyboard, with your wheel mouse as well. If you copy the configuration file to `~/.imwheelrc`, then you can also make changes as a normal user, not just as 'root'.

Left-handed usage of the mouse To swap the left and right mouse button functions, you can enter the following command:

```
newbie@earth:/home/newbie > xmodmap -e "pointer = 3 2 1"
```

You can also include this command in the file `~/.xinitrc`.

Two mice or a touch pad It is no problem if you want to include a second mouse or a touch pad. Again, we would advise you to use SaX2. Chose the 'Custom' configuration mode and in the 'mouse' menu, add a further mouse symbol, using the 'New' button. This second mouse will not be configured automatically, but you can enter the data yourself by hand. Under 'Drivers' you can chose the relevant driver `mouse` if this is another mouse, or you can select one of the drivers listed for touch pads. Under 'Protocol' and 'Connection', the procedure is similar. The setting 'Protocol' only makes sense if you are using the `mouse` driver. Otherwise you should select the pads type settings in the Expert dialog. If you have a touch pad driver, select `None` in the 'Protocol' section.

In the 'Expert' menu then choose additional options in 'Miscellaneous'. The option `Send Core Events` deserves special mention here; if you activate this action you can use input devices in a parallel fashion. Otherwise you can switch between the two input devices using the command `xsetpointer`. You can use the command

```
newbie@earth:/home/newbie > xsetpointer -l
```

to have a list of possible input devices displayed.

Keyboard

By nature there is not as much to write about the subject of "Keyboard" as there is on other input devices. Usually there is nothing special to configure here.

The following programs and files can however be useful in certain cases: To make changes or just to test the keyboard layout during operation there is the program `xkeycaps`, package `xkeycaps`, series `xap`. In the file `~/.Xmodmap` you can then save changes permanently to individually changed characters. For the format of this file, see the manpage for `xmodmap` (`man xmodmap`).

In KDE the program `kikbd` (“International keyboard layout”) is available in the ‘System’ menu of KDE. With this you can switch very conveniently between different keyboard layouts. This program is very useful if, for example, you want to use an English keyboard for programming (here the brackets `{}` and `[]` are much more easily accessible) and a German keyboard to write texts.

6.6.3 3D Acceleration

For information on 3D acceleration, please refer to the “Configuration” manual, in the Chapter “OpenGL/3D Configuration”.

7 Linux and Hardware

7.1 Preliminary Notes

It is now possible to integrate almost all hardware components into a Linux system with little (or more) trouble. How this is achieved in individual cases and which software is available is the subject of this chapter. For problems involving *printing*, see Chapter 9 page 197. *PCMCIA hardware* is described in Chapter 8 page 183.

7.2 Cards

In a standard PC, normally the PCI (“Peripheral Component Interconnect”) bus is used. Most times the ISA (“Industry Standard Architecture”) bus is still present in systems.

The AGP (“Accelerated Graphics Port”) is an exception to this, as devices on this port are also visible on the PCI bus. They are treated logically in exactly the same way as PCI bus cards. PCMCIA bus cards are found mostly in portable computers such as laptops and notebooks. These cards are described in Chapter 8 page 183.

7.2.1 PCI and ISA Cards

We shall now take a closer look at PCI and ISA bus cards. The configuration of these cards will then be explained using sound cards as an example.

PCI cards

PCI bus cards normally don’t need to be configured at all by the user. The agreement on the PCI bus standard was an attempt to remove many restrictions of the old bus system. In this respect, a sensible automatic configuration of the cards was planned. Each PCI card is activated by the computer’s BIOS when booting. In many cases you can influence the distribution of the interrupts through settings in the computer BIOS. When Linux starts, it reads the configuration of the PCI devices directly from the PCI BIOS and, from this point on, uses this data for all information about the PCI subsystem.

With the command

```
earth:/ # lspci -tv
```

you can be shown a list of all the devices on the PCI bus recognized by Linux; see output 7.2.1.

```
-[00]--00.0 Intel Corporation 440BX/ZX - 82443BX/ZX Host bridge
+-01.0-[01]----00.0 Nvidia Corporation Riva TNT
+-04.0 Intel Corporation 82371AB PIIX4 ISA
+-04.1 Intel Corporation 82371AB PIIX4 IDE
+-04.2 Intel Corporation 82371AB PIIX4 USB
+-04.3 Intel Corporation 82371AB PIIX4 ACPI
+-06.0 Adaptec 7890
09.0 Digital Equipment Corporation DECchip 21140 [FasterNet]
```

Output 7.2.1: Output of the Command `lspci -tv`

In Subsection 7.3 page 169 we will show you how to configure a PCI sound card.

ISA cards

The ISA card is the oldest form of PC-compatible card. The bus is 8 or 16 bits wide, and has a maximum frequency of 8 MHz. Many motherboards provide the option of increasing the frequency of the ISA bus — but this often leads to problems. ISA bus cards can be divided into three categories:

Legacy cards: These cards must be configured entirely by hand. Resources such as I/O addresses, interrupts and DMA channels must be set through jumpers or microswitches on the cards themselves. You must also ensure that resources such as IRQ or DMA channels are not used by two or even more cards simultaneously. Therefore the configuration of such cards often turns out to be somewhat difficult.

Jumperless cards: Jumperless cards, compared to the first generation of ISA cards, have the advantage that the setting of resources is no longer made via jumpers, but via a special configuration program. These programs, however, are usually only available for DOS, and thus cannot be used in Linux.

PnP cards: Plug and Play cards are the logical progression from jumperless cards. These cards contain data on their configuration and a list of possible configurations. A special program or a driver can now query the PnP cards in the system and configure all cards so they do not conflict with each other. If a card has not been configured, then in practice it is not available to the system.

To configure PnP cards in Linux running a 2.2.x kernel, the package `is-
apnp` from the series `ap` is used. This package contains the two programs, `pnpdump` and `isapnp`.

How are PnP cards activated with `isapnp` tools?

Proceed as follows:

- Change to the user `'root'`.

- If the file `isapnp.conf` already exists in the directory `/etc`, you should first make a backup copy by changing its name.

```
earth:/ # mv /etc/isapnp.conf /etc/isapnp.conf.bak
```

- By entering

```
earth:/ # pnpdump -c > /etc/isapnp.conf,
```

you will create the file `/etc/isapnp.conf`. The parameter `-c` ensures that `pnpdump` prepares the output file in such a way that the PnP devices are activated immediately.

- If you enter

```
earth:/ # isapnp /etc/isapnp.conf,
```

the PnP cards in your computer will be activated. — In SuSE Linux this is activated automatically each time the system is booted.

Possible sources of error

Symptom: `pnpdump` gives the output: "No boards found". It may also be the case that one or more of the installed ISA PnP cards cannot be recognized.

Possible Explanations:

- There are no PnP cards installed on your system: check which cards are installed on your system and read the documentation about them. If necessary, ask the dealer from whom you bought your computer which cards are installed on it.
- The cards are defective: check if the cards are correctly positioned in the slots. See if the cards work correctly under a different operating system.
- There are certain cards which can either be run as ISA PnP cards, or which can be allocated fixed resources. (Example: a number of 10 MBit NE2000 ISA Ethernet cards). Usually you can switch between the two modes using a DOS program.

Depending on the configuration of such cards, these do not appear in the output of `pnpdump`.

- There are a few ISA PnP cards which will not work together with other ISA PnP cards in the same system. In such cases, only one of the cards is visible in the output of `pnpdump`. It is possible that one of the cards can be configured so that it is no longer driven as an ISA PnP card (for example, with a jumper directly on the card so that it uses fixed resources).
- There seem to be (in very few cases) cards which perform a reset after `pnpdump` is run. If this involves a SCSI card, for example, a controller on which the root partition is dependent, this can cause the computer to hang. If this problem occurs, you must try and configure the card in such a way that it no longer functions as an ISA PnP card.

Otherwise you can do away with initializing the card in Linux completely and boot the system with the help of `loadlin`; in this case, the card is already initialized.

Symptom: When running the command `isapnp` or when booting, you will see an error message like this:

```
* LD setting failed, this may not be a problem.
* Try adding (VERIFYLD N) to the top of your script
*
* Error occurred requested 'LD2' on or around line 319
* --- further action aborted
```

Solution: Follow the advice of the error message, and at the beginning of the file `/etc/isapnp.conf`, insert the line

```
(VERIFYLD N)
```


that is:

```
# [...]
# (DEBUG)
(VERIFYLD N)
(READPORT 0x0203)
(ISOLATE)
(IDENTIFY *)
# [...]
```

Symptom: When running the command `isapnp` or when booting, you see an error message like this:

```
[...]
/etc/isapnp.conf:66 --
  Fatal - resource conflict allocating 16 bytes
  of IO at 220 (see /etc/isapnp.conf)
/etc/isapnp.conf:66 -- Fatal - IO range check
  attempted while device activated
/etc/isapnp.conf:66 -- Fatal - Error occurred
  executing request '<IORESCHECK>' --- further
  action aborted
```

Solution: You probably have a conflict between the values selected in the file `/etc/isapnp.conf` and those resources already being used in your system. You should compare your `/etc/isapnp.conf` with the information on the resources in question, which can be found in `/proc`; this should not be necessary if you allow `pnpdump` to undertake this allocation itself by using the option `-c`.



Caution

In many cases difficulties can arise if, in the setup of the computer, the automatic configuration of ISA PnP resources is switched on through the BIOS. You should switch this off. Since the configuration is performed by means of `pnpdump` or `isapnp`, this setup option is not necessary.

In case of problems which go deeper than this, you should consult the extensive documentation on the package `isapnp`. This can be found on your system in the

directory `/usr/share/doc/packages/isapnp`. At the end of this subsection we will take a closer look at configuring a PnP card, using the example of a sound card.

7.3 Sound Cards

Today, mainly PCI sound cards are available on the market. Sometimes you may also find ISA bus sound cards.

In SuSE Linux there are now two methods for supporting sound cards. The simplest method is to use YaST2. YaST2 uses drivers which were developed by the ALSA project. Information on this project can be found at the URL:

<http://www.alsa-project.org>. You also can install the ALSA drivers with the program `alsaconf`.

7.3.1 Configuring Sound Cards with YaST2

To configure your sound card (if you didn't do this during the installation, e.g. because you added the sound card later to the system) just start YaST2 as user 'root' and choose the module 'Hardware/Sound'.

YaST2 now tries to automatically detect your sound card and brings up a list of drivers and corresponding sound card models suitable for your hardware. Select the driver and sound card model matching your card. Consult the sound card manual if necessary. If the sound card was not autodetected, click on 'Manual' and choose your sound card and the corresponding driver from the list. By selecting 'Options', you can enter some options specific for the selected card. The manual to your sound card gives more details.

By selecting 'Next' you will get to the "Test screen" to check the configuration of the sound card. The buttons '+' and '-' allow you to change the sound volume. You should start at a level of about 10% to avoid damage of your audio system or your ear. Also check the volume control of your audio system. By clicking on 'Test' you should hear a test sound. If there's no sound, you should check the volume of your audio system first. If you heard the sound and the volume is ok, click on 'Next'.

This finishes the configuration and the volume settings are saved. Click on 'Next' again to get back to 'YaST2 System configuration and administration'.

You can change the volume and mixer settings later by using programs like `kmix` or `gamix`.

7.3.2 Configuring Your Sound Card With `alsaconf`

ALSA (Advanced Linux Sound Architecture) is supported by many developers and the number of drivers available for it is growing steadily. Along with a native interface, ALSA offers the `OpenSoundSystem`, an interface compatible to `OSS`. Thus you can continue to use applications that were originally developed for

OSS. All the sound configuration options can only be taken advantage of fully via the ALSA interface itself.

ALSA is completely based on the modular concept of the kernel. The required modules are located in the kernel module directory (`/lib/modules/...`) in the `misc` directory. Configuring ALSA takes place by means of entries in the `/etc/modules.conf` file. YaST2 takes care of creating these entries for you. However, ALSA offers its own program as well, `alsaconf`, which can be used in place of YaST2 for configuration.

With ALSA, you can run multiple sound cards simultaneously at any time. To do this, the `/etc/modules.conf` file only needs to be extended using the required parameters.

Proceed as follows to integrate your sound card into your SuSE Linux using `alsaconf`:

- Log in to a text console as user `root`.
- Start the program `alsaconf`

```
earth:/ # alsaconf
```
- `alsaconf` will at first attempt to detect the sound card. It not only searches for ISA bus cards, but also for PCI bus cards. Plug and Play cards on the ISA bus are scanned as well.
- `alsaconf` shows you a list of the sound cards it found and offers a configuration for them. If a card has not been found, a selection list of potential cards will immediately appear.
- Now you will be asked whether the `/etc/modules.conf` file should be modified. You should confirm this by selecting `YES`.
- The necessary entries are now created in `/etc/modules.conf` and the driver will be given a trial run.
- Subsequently, you have the option of testing the sound card. A “test sound” will be played back for this purpose. Do not turn the speakers up too loud.
- The drivers will automatically be loaded the next time you start your system. You can also manually load and unload them.

Enter

```
earth:/ # rcalsasound start
```

to load and

```
earth:/ # rcalsasound stop
```

to unload the drivers.

7.4 Ports on a Computer

7.4.1 PS/2 Ports

PS/2 ports were developed by IBM for the PS/2 platform. The ATX Standard, in which this type of port is included, has found widespread use in the last few years. You can attach a mouse and a keyboard on the PS/2 port.

Configuring the PS/2 port is not necessary under Linux because it is supported by the kernel. The system automatically detects existing PS/2 ports and can control the devices attached to them. The ports can be found and addressed under the device name `/dev/kbd` and `/dev/psaux`.

7.4.2 Serial Ports

The serial port (RS232) of your computer is often used for external modems. You can also attach a mouse with an RS232 plug on this port.

Support for serial ports is provided by the kernel module `serial.o`. This module is automatically loaded when the port is addressed by certain software. For instance, supposing you've configured your Internet access such that a modem on COM1 is used, the module is automatically post-installed before the connection to your provider is established.

You can change your serial port configuration using the program `setserial`. This program is run at every system start by the script `/sbin/init.d/serial`. For more information on `setserial` see the man page for this program. Serial ports can be found and addressed under the device file(s) `/dev/ttyS0`, `/dev/ttyS1`, etc.

7.4.3 Parallel Ports

The parallel port of your computer is primarily used to connect printers. Many other devices such as ZIP drives and CD-ROM/RAM drives can also be connected to this port.

The “parport” subsystem of the current Linux kernels 2.2.xx is designed to operate a number of devices on the parallel port *simultaneously*.¹ As an example, you can attach a ZIP drive to the parallel port and, in turn, connect a printer to the ZIP drive.

Initializing `parport`

The `parport` and `parport_pc` kernel support are required for this. `parport` is the actual subsystem of the kernel and `parport_pc` takes care of the hardware connections of the PC port as well as for those of some AXP computers. With the command

¹For the “old” kernel 2.0.xx only one device can be used *exclusively* on a port; as a rule, “multiple” connections also cannot be used!

```
earth: # modprobe parport_pc
```

you can attach the modules manually. Check if the ports could be found and configured by entering the command `lsmod`; the `parport` and `parport_pc` ought to appear in the list of modules. You can also look in the file `/var/log/messages` to check which values were used to initialize `parport` (see File contents 7.4.1).

```
Jun  3 09:15:53 tux kernel: parport0: PC-style at 0x378 [SPP,
                                         ECP]
```

File 7.4.1: `/var/log/messages`: `parport` Initialization

The entries in the file 7.4.2 are responsible for configuring the interfaces; the entry which is commented out indicates how to integrate another parallel interface.

```
alias parport_lowlevel    parport_pc
options parport_pc io=0x378 irq=None

# If you have multiple parallel ports, specify them this way:
# options parport_pc io=0x378,0x278 irq=None,none
```

File 7.4.2: `/etc/modules.conf`: `parport` Configuration



Note

In case of problems, you should first check in the BIOS of the computer how the port in question is specified there. The port should have its own explicit address (if possible, `0x378`) and its own interrupt. “Automatic” settings should be avoided.

Sometimes it is also necessary to switch on the devices in the correct sequence; please look in the appropriate handbooks! One recommended method is to first switch on the computer, and then – before booting has commenced – to *immediately* switch on the external devices.

Using `parport`: Printers, ZIP Drives, PLIP, etc.

If one parallel port device is to be used, its relevant module must also be loaded; in the case of a printer, this is the `lp` module:

```
earth: # modprobe lp
```

`lsmod` ought to show us, amongst other things:

As with other cases, further details are supplied in the file `/var/log/messages` and also in the `proc` file system in `/proc/sys/dev/parport/`. If everything is running correctly, you won’t need to load the modules manually – `Kmod` will


```

parport_pc          5568    1  (autoclean)
parport             6884    1  [parport_pc lp]
lp                  5116    0  (unused)

```

do this for you automatically (Section 10.2 page 231) if, for example, a printer job is sent. ZIP drives are not served either by `ppa` or by `imm` — see below, Section 7.5.3 page 177.

Further Suggestions

Look in the kernel sources, in directory `/usr/src/linux/Documentation`, at the files `parport.txt` and `paride.txt`.

Information on the current status of `parport` programming can be found on the web at <http://www.torque.net/linux-pp.html> and <http://www.torque.net/parport/>.

7.4.4 USB – Universal Serial Bus

In 1994 Compaq, Intel, Microsoft and NEC cooperated to setup the specifications for the “Universal Serial Bus” (USB). They aimed to create an easy-to-use connection to the telephone and an extension of the already available connection types (serial, parallel). Using this bus system, you are able to connect various devices like mice, keyboards, printers, scanners and joysticks to your computer. The different cabling used to connect these devices should be replaced by one , unified connection type. The USB bus is capable of mounting and unmounting devices to your computer while it is running – the computer does not need to be rebooted.

The first version of the USB bus was released in 1996 and extended it’s capabilities from release to release. In the current version (USB 1.1) 127 devices are supported. The maximum bandwidth is limited to 12 Mbits/s.

In Linux, support for the USB bus is constantly growing. Numerous devices are already supported and can be configured easily. From kernel version 2.4 on, the usage of USB devices with Linux will be standard.

Available Modules

In the following section we are going to give you an overview on the available modules and will give you details which module should be loaded for which device. The listed modules are located in the directory `/lib/modules/2.4.9/usb`. Details on the modules can be found in the directory `/usr/src/linux/Documentation/usb/`.

If we talk about “input events” or “events” in general here, you should think about mouse movements, keystrokes, etc.

input.o This module needs to be loaded before other USB input devices can be supported by Linux. It is used to control the communication between the hardware (via USB) and the processing of the events. More information on

that can be found in the file `input.txt` in the directory `/usr/src/linux/Documentation/usb`.

usbcore.o USB core is the USB-subsystem that simplifies the communication with the installed hardware which depends on the data structures, macros and functions of the hardware. By using a general interface to the operating system (API) common routines of all USB devices and host controllers are supported.

usb-uhci.o / usb-ohci.o This module enables the USB host controller of your motherboard. Most of the motherboards available today have the USB host controller integrated in the chipset. Older motherboards can be upgraded by a PCI card. The USB host controller is then integrated on this PCI card. There are two standards making the USB host controllers compatible: The Universal Host Controller Interface (UHCI) Intel and the Open Host Controller Interface (OHCI) defined by Compaq.

The module `usb-ohci.o` enables the OHCI-chipset (Compaq, NEC, iMacs, OPTi, SiS, ALi, Lucent, CMD, etc.) whereas the `usb-uhci.o` module enables the UHCI-chipset (Intel, VIA) controllers.

Tip

If you are logged in as user `root`, you are able to read out the type of USB host controller on your motherboard by using the command `lspci -v | grep USB`. You can then load the respective module.



hid.o HID (Human Interface Devices) is a device class accessed by this generic module. Supported devices are keyboards, mice, joysticks, digitizers etc.

For some devices there are two modules available in the modules directory: The `*dev.o` module is a “high-layer”-module (“*” standing for e.g. “mouse”) and is meant to as an interface between the hardware and the users program like “X” or “games”; `usb*.o` is a “low-layer”-module which generates the input events. Both modules are needed, whereas the “low-layer”-module can be replaced by the `hid.o` module.

There are two modes to describe the event input of a USB device: *HID* enables support of all functions of a device so requires the quite complex module `hid.o`: *HIDBP* (HID boot protocol) is a very simple mode and only supports the basic functionality of a device to make it accessible by the computers BIOS.

usbkbd.o This module enables the usage of the basic functions of a keyboard in the *HIDBP*-mode, but for this device class the access should be enabled via the `hid.o` module. Unlike the `hid.o` module, only keyboards will be enabled with this module.

usbmouse.o This module enables using the basic functionality of a mouse in *HIDBP*-mode. The usage of the `hid.o` should be preferred as this enables event management for all other input devices as well.

keydev.o This module “translates” input events to the architecture dependent RAW mode of the keyboard. This module is still under development but works fine for almost every architecture.

mousedev.o This module converts the input events of a mouse, digitizer or tablet to PS/2 compatible events. As it is not finished yet, you should use the `evdev.o` module or, if you are using a “wheel” mouse, the `hid.o` module.

acm.o Load this module if you plan to use communication devices like modems or ISDN adapters of the device class “ACM” (Abstract Control Model). “ACM” is a method to control modems. You can find a list of supported modems in the file `acm.txt` in the directory `/usr/src/linux/Documentation/usb`.

audio.o enables audio devices and supports digital mixers, wave devices and loudspeakers.

dabusb.o supports amateur radio devices.

dc2xx.p enables digital cameras of the *Kodak DC-2xx* series and also some non-*Kodak*-cameras. You can find more information on this in the file `dc2xx.txt` in the directory `/usr/src/linux/Documentation/usb`.

evdev.o “evdev” is a general interface (API) controlling the input of devices as event codes are identical and hardware-independent on all architectures. This module supports this API.

ibmcam.o is used for supporting the *IBM(C-It)* and *Xirlink-Webcams*, but not all features are supported now. You may find more details on that in the file `ibmcam.txt` in the directory `/usr/src/linux/Documentation/usb`.

joydev.o is used to communicate with joysticks.

ov511.o enables the “OV511” chip from *Omnivision Techonolgies* which is used in some webcams. In addition, the *OV511/OV511+* and *OV7610/20/20AE CCD*-chips are supported. This driver is compatible with most “Video4Linux” applications, but not to some video conferencing applications. For details on the supported resolutions, please read `ov511.txt` in the directory `/usr/src/linux/Documentation/usb`.

printer.o is used to communicate with USB printers.

scanner.o was developed to support USB scanners from *Hewlett Packard* (HP) first, but became more and more hardware independant later. Read more about it in `input.txt` in the directory `/usr/src/linux/Documentation/usb`.

usb-serial.o With this module you can control serial devices on the USB bus. The serial signal are converted to USB signals. Up to 256 connections can be used simultaneously. Details can be found in the file `usb-serial.txt` in the directory `/usr/src/linux/Documentation/usb`.

usb-storage.o enables USB mass storage devices like ZIP, floppy disk drives, hard disks and smart media card devices. The driver isn't quite stable at the moment.

wacom.o enables support for the *Wacom Graphire and Intuos Tablets*. Although support should already be included with the *HID* module, it was only possible to use a *Wacom PenPartner*.

wmforce.o enables the use of *Logitech WingMan Force* joysticks with USB interface. The driver works well, but the "Force Feedback" function is not yet supported.

Configuring a USB mouse and keyboard

For a common configuration with USB mouse and keyboard the following modules should be loaded:

- `input.o`
- `mousedev.o`
- `keybdev.o`
- `usbcore.o`
- `usb-uhci.o` or `usb-ohci.o`
- `hid.o`

Further sources

More information on the USB implementation and on supported devices can be found at

- <http://www.linux-usb.org>
- <http://www.qbik.ch/usb/devices/>

7.5 Removable Drives

A variety of removable drives can be used under Linux: Floppy disk drives, ZIP, JAZ or SyQuest Drives. Magneto-optical drives can also be used in Linux.

7.5.1 Floppy Disk Drives

The "mtools" were developed for easy access of MS-DOS formatted floppy disks.

You can also read and write to disks formatted with non-MS-DOS file systems (e.g. `ext2` or `minix`). Many other file systems are also available. In order to accomplish this, however, you'll first need to use the command `mount` to access the medium.

7.5.2 LS-120 Drives

LS-120 drives are connected at the (E)IDE port of the computer and are seen by the system as hard disks, meaning that they, too, need to be accessed first with the command `mount`.

7.5.3 ZIP Drives

There are ZIP drives for different interfaces: parallel, ATAPI and SCSI, as well as USB.

For IDE and SCSI, you don't need a special device driver. They can be connected on the bus and addressed without extra configuration using the bus's device driver (IDE or SCSI). The devices are addressed by the device names,

```
/dev/hda - /dev/hdd for IDE (ATAPI)
```

and

```
/dev/sda - /dev/sdm for SCSI
```

The device needs to be tied into the directory tree with the command `mount`.



Tip

ZIP disks often come preformatted with an MS-DOS file system; usually you'll be able to address such a partition as `/dev/hda4` or `/dev/sda4`. More information about this is included in file:///usr/share/doc/sdb/de/html/maddin_zip.html.

The installation of the *parallel port* version proves to be somewhat more complicated. You'll need SCSI hard disk support, `parport`, `parport_pc` support (see Section 7.4.3 page 171) as well as `ppa` or the `imm` driver of the kernel.² `imm` is needed for more modern drives (e. g. for the ZIP 250). Furthermore you should make sure that, in the BIOS of the computer, the EPP mode for the parallel port is set. If you are unsure as to which kind of drive you have, try loading `imm`:

```
earth: # modprobe imm
```

If this doesn't work the same with `ppa`:

```
earth: # modprobe ppa
```

With this, the `parport` subsystem will be automatically initialized. If this does not function, look at Section 9 page 171.

Sometimes an entry in `/etc/modules.conf` can help matters (see File contents 7.5.1 on the following page). The alias `scsi_hostadapter` may only be used if you don't have a standard SCSI host adapter in your system. An alternative is to write the necessary `modprobe` commands into the boot script

²The `ppa` and/or the `imm` driver are contained in the "SCSI low-level drivers" – but you *don't* need to compile your own kernel; the necessary modules are precompiled and included with SuSE Linux

/sbin/init.d/boot.local; For more information on this file, see Section 14.4 page 309.

```
# alias scsi_hostadapter ppa
pre-install ppa modprobe "-k" parport_pc
```

File 7.5.1: /etc/modules.conf: ppa Configuration

After this, ZIP disks can be accessed in the same manner as SCSI hard disks. It is also necessary to “mount” the medium.

Adding support for USB ZIP drives is quite easy to manage. First, make sure that the USB subsystem is loaded (see Section 7.4.4 page 173). The following modules need to be loaded:

```
usbcore
usb-uhci
```

or

```
usbcore
usb-ohci
```

In addition you have to load the module named `usb-storage`.

Then you can access the ZIP drive just like SCSI hard drive. You need to mount the media first.

7.6 Modems

7.6.1 External Modems

External modems are normally attached directly to the serial port of the computer. Programs can address them via device files with the name(s) `/dev/ttyS0`, `/dev/ttyS1`, etc.

7.6.2 Internal Modems

Internal modems in desktop machines are either PCI or ISA devices. Please see Section 7.2 page 165 for information about these devices.

Since modems are accessed through serial ports, internal modems must make these ports available for themselves.

Table 7.1 gives you a list of standard resources for serial ports.

PCI modems

If you have a PCI modem, it is most likely a “Winmodem”. Winmodems are not modems. They are hardware devices which need a special driver that allows them to emulate a real modem. In general, the driver is provided by the maker, and unfortunately there are almost no such drivers currently available for Linux.

Device file	Name	Interrupt	IO-Address
/dev/ttyS0	COM1	4	0x3f8
/dev/ttyS1	COM2	3	0x2f8
/dev/ttyS2	COM3	4	0x3e8
/dev/ttyS3	COM4	3	0x2e8

Table 7.1: Standard Resources for Serial Ports

You can obtain more information about Winmodems and how they are supported under Linux in the SDB article file:///usr/share/doc/sdb/en/html/cep_winmodem.html.

If the modem is not a Winmodem, you should be able to use it in SuSE Linux. A certain amount of manual work is necessary, however, since these modems cannot be configured automatically.

ISA modems

With ISA modems, there are two main considerations:

- The modem itself needs to be configured (IRQ and IO addresses for the port).
- The kernel needs to know what IRQ and port the card is using.

Initializing the modem

As described in Section 7.2 page 165, ISA cards can be jumpered, jumperless or Plug and Play. The initializing of PNP cards is also described in that section.

Some suggestions about settings:

Setting your modem to use /dev/ttyS1 (COM2) with IRQ 3 will usually present the least number of problems. /dev/ttyS3 (IRQ 4) can also be used, but you may have a conflict if you are also using ttyS0 (COM1), as this port also normally uses IRQ 4. So if you have the mouse on /dev/ttyS0 you should not use /dev/ttyS0 or /dev/ttyS2.

Generally, IRQ 5 and IRQ 7 should also be avoided, as these are often used for sound cards or printers. In fact, if you are going to use a sound card, you should configure it first because it claims many resources.

Passing parameters to the kernel

After you have established the port and IRQ for your modem, you need to tell the kernel what values the card is using. If you have used a standard port and IRQ (/dev/ttyS0 or /dev/ttyS1), your card will be automatically recognized the next time you reboot.

If you are using other IO port/IRQ combinations, you can manually tell the kernel the device's IRQ, using `setserial`:

```
earth:/ # setserial /dev/ttyS3 irq 10
```

You can now configure the modem for Internet connection.

If the modem is recognized, you can automate the `setserial` command by modifying the appropriate line in the file `/sbin/init.d/serial`:

Change the line:

```
# run_setserial /dev/ttyS3 $AUTO_IRQ autoconfig
```

to:

```
run_setserial /dev/ttyS3 irq 10
```

7.7 Tape Drives

Tape drives, also known as streamers, are mainly used for high-capacity backup purposes. There are tape drives on the market for a variety of ports.

SCSI Tape Drives

Tape drives that attach to the SCSI bus are the most common. You can assume that all devices that are attached to a Linux-supported SCSI controller can be used, although if your device has extra functions, special software is needed to access these.

A SCSI tape drive is controlled in the following way:

- The first tape drive in the system is controlled using the device files `/dev/st0` and `/dev/nst0`, the next with `/dev/st1` and `/dev/nst1`, etc. `/dev/stX` stands for a drive that automatically rewinds after use. `/dev/nstX` on the other hand stands for a tape drive that leaves the tape in the current position after use.

We suggest using `/dev/nstX`, because, using `mt`, you have exact control of the drive.

- Use the command `mt` to control a tape drive:

```
earth:/ # mt -f /dev/nst0 retension
```

Adjusts the tape tension by rewinding and then fast-forwarding the tape.

```
earth:/ # mt -f /dev/nst0 rewind
```

Moves the tape to a point just after the last recorded information. Used to add information to a partially used tape.

```
earth:/ # mt -f /dev/nst0 eof
```

Moves the tape to a point just after the last recorded information. Used to add information to a partially used tape.

```
earth:/ # mt -f /dev/nst0 erase
```

Erases all information on a tape. Be careful, there is no way to recover erased information!

IDE Tape Drives

IDE tape drives are connected to the IDE bus of the computer. The drive is found automatically by Linux. You can address the device under the device files `/dev/ht0` or `/dev/nht0`.

You can control the device using the command `mt`, as given in Section 7.7 on the facing page.

Floppy Tape Drives

Floppy tape drives are simply attached to the floppy disk port of your computer (almost every computer has a floppy port).

Because these devices need to be controlled with precise timing, errors often occur while saving and writing backups.

To use these devices in Linux, you'll need a special driver, the `ftape` driver. This driver is included in the kernel. Read the documentation in `/usr/src/linux/Documentation/ftape.txt` for more information on the configuration of this device. The entire documentation for the `ftape` subsystem can be found on the device driver developers' home page: <http://www.instmath.rwth-aachen.de/~heine/ftape/>.

8 Notebooks — PCMCIA, APM, IrDA

Notebooks, in particular, have special components and requirements, among these are “Advanced Power Management” (APM), Infrared ports (IrDA) and PC cards (PCMCIA). Occasionally these components can also be found in desktop computers. Because the differences between the two types are insignificant, the use and configuration of both will be described in this chapter.

Those interested in specific notebooks should visit the “Linux laptop home page” at <http://www.cs.utexas.edu/users/kharker/linux-laptop>. Another good source of information is the “Moblif” home page at <http://mobilix.org/index.html> (*Mobilif — Mobile Computers and Unix*). There you can find an interesting Laptop HOWTO as well as an IrDA HOWTO. In addition, there is an article *Laptops and Notebooks (PCMCIA) in Linux* (`file:/usr/share/doc/sdb/de/html/laptop.html`) in the SuSE support database.

8.1 PCMCIA

8.1.1 Hardware

PCMCIA stands for “Personal Computer Memory Card International Association” although it is commonly used as a collective term for both the hardware and the associated software. The essential part is the PCMCIA card, of which two types exist:

PC cards: This is the most common type. With a 16-bit bus bandwidth, most are relatively inexpensive and are usually trouble-free and stable.

CardBus cards: This is a new standard. They have a 32-bit bus bandwidth and are quicker, although also more expensive. Because the data transfer rate is often limited at other points in the system, the increase in bandwidth is often unjustified. In the meantime there are quite a few drivers for these cards as well, although they are often unstable — depending on the PCMCIA controller.

When the PCMCIA service is active, the card type can be found with the command `cardctl ident`. A list of supported cards can be found in `SUPPORTED_CARDS` in `/usr/share/doc/packages/pcmcia`. There you’ll also find the actualized version of the `PCMCIA-HOWTO`.

The second important component is the PCMCIA controller, or the PC card or CardBus bridge. This creates a connection between the card itself and the PCI

bus or, in older devices, the ISA bus. These controllers are most often compatible to the `i82365` chip from Intel; all models are supported. The type of controller can be found with the command `probe`. In the case of PCI devices, the command `lspci -vt` also offers interesting information.

8.1.2 Software

All necessary drivers and programs can be found, as long as they are not already integrated into the kernel, in package `PCMCIA`, series `a1`. The modules `pcmcia_core`, `i82365` (or `tcic`, less often) and `ds` form the basis and are normally automatically started when booting. They initialize the PCMCIA controller and supply basic functions.

Because PCMCIA cards can be added or removed at run time, a daemon is needed to check the activities of the slots. This is done by the Cardmanager (`cardmgr`), which is automatically started after loading the base modules. When a card is inserted into a slot, Cardmanager determines the type and function and loads the necessary module. You can check which modules are loaded with the command `lsmod`. If all modules are successfully loaded Cardmanager, depending on the function of the card, starts predetermined initialization scripts that build the necessary network connections or *mount* the necessary external SCSI drives. When the card is removed Cardmanager, using the same scripts, ends the diverse card activities, after which the modules that are no longer needed are unloaded. Theoretically, you can simply remove the card. This works very well with network, modem, or ISDN cards as long as there is no active network connection. It does not work with mounted partitions, external drives, or NFS directories. For this you need to be certain that the devices are synchronized and cleanly *unmounted*. If in doubt the following command can be of help

```
earth: # cardctl eject
```

This command deactivates the card, as long as it is still in the notebook.

8.1.3 Configuration

When the PCMCIA package is installed, four PCMCIA variables will be found in `/etc/rc.config`. `START_PCMCIA` determines whether the service should be started when booting. It can be started with the command `rcpcmcia start` even when `START_PCMCIA` is set to `no`. `PCMCIA` determines the type of PCMCIA controller. This value is automatically set at installation and almost always has the value `i82365`. The other two variables, `PCMCIA_PCIC_OPTS` and `PCMCIA_CORE_OPTS` involve options for the basis modules and can normally be left empty.

Because the selection of driver modules is taken care of by Cardmanager (`cardmgr`) no other hardware related settings are necessary.

Ethernet and Token Ring

Ethernet or token ring connections can be easily configured using YaST. In the menu item 'System administration' → 'Network configuration',

these devices can be configured in the same manner as normal hardware devices with the exception that the device must be marked as a PCMCIA device with **(F9)**; see also Section 3.7.3 page 74. The following points should be taken into consideration:

- When more than one ethernet or token ring device is marked as PCMCIA, only the first device configuration is used.
- The settings, in contrast to normal network cards, are immediately activated with **(F10)**.
- The device names shown (`eth0`, `eth1`, `tr0...`) should only be seen as information relating to the type of device because the numeration of PCMCIA devices is determined dynamically.
- When a non-active device is reactivated with **(F4)** the device must also be marked as a PCMCIA device with **(F9)** again.

ISDN

ISDN-PC cards are configured in much the same manner as other cards. The only difference is that the device must be marked as a PCMCIA device under the menu items ‘System administration’ → ‘Network configuration’ with **(F9)**. In the ‘Configuration of ISDN hardware’ dialog, IRQ, IO port, and the ISDN protocol (Euro-ISDN or 1TR6) should not be changed.

ISDN modems are also available as PCMCIA cards. They are modem or multi-functional cards with an additional “ISDN Connection Kit” and are handled as modems.

Modems

Modem PC cards normally have no specific settings. After insertion, the card is available in `/dev/modem`. The configuration of the device with `wvdial` (**wvdial**) follows that of a normal modem.

SCSI and IDE

The necessary driver modules are loaded by Cardmanager. After insertion, the device is ready for use. The device name is determined dynamically. Information about available SCSI or IDE devices can be found under `/proc/scsi` or `/proc/ide`.



Note

Before inserting external hard disks, CD-ROM drives and similar devices be sure that the device is on. SCSI devices must be actively terminated.

Important: Before a SCSI or IDE device is removed from the slot, the partitions or device must be unmounted. If this step is left out, the device can only be accessed after rebooting although the system is still stable.

It is possible to install Linux entirely on such an external drive, but the boot process is somewhat more complicated. A “boot disk” will be needed which includes the kernel and an init RAM disk (`initrd`); more information on this can be found in Section 13.4 page 288. The `initrd` includes a virtual file system with all the necessary PCMCIA modules and programs. The SuSE Linux “boot disk” and boot disk images include these, enabling booting of your external installation. It is, however, somewhat uncomfortable having to load the PCMCIA support every time you boot your system; advanced users can create their own boot disks that are tailored to their systems¹.

8.1.4 Configurations for Changing — “Schemes”

Mobile computers often have different configurations. For example, one for work and one for home. For PCMCIA devices this proves to be quite easy, although the configuration files need to be edited by hand. Note that the following description does not apply to ISDN and modem cards. Please also read `/usr/share/doc/packages/pcmcia/README.SuSE` where you can find the most up-to-date information as well as some tips and tricks.

Singular configuration profiles for PCMCIA are known as “schemes”. The configuration files found in `/etc/pcmcia/*.opts` can include more than one scheme. The data for a specific scheme is included in configuration blocks and selected with an “address”. The addresses are comma-separated words. The first word is the name of the scheme. The other words contain information such as the slot in which the card is inserted or the number of a hard disk partition. Detailed information can be found at the beginning of the `/etc/pcmcia/*.opts` files as well as in the PCMCIA HOWTO. The standard scheme used by SuSE Linux has the name `SuSE`. The `SuSE` scheme is edited with `SuSEconfig`, when the network connection is configured with `YaST`. Due to this, manual alterations to this scheme will be lost by using `SuSEconfig`.

We’ll use an ethernet card configuration as an example scheme configuration, making two schemes named `work` and `home`. The file `/etc/pcmcia/network.opts` (see file 8.1.1 on the next page) will have several configuration blocks with the following addresses:

- `SuSE, *, *, *`: This block already exists and should be left as is, in case you’d like to use `YaST` for future configurations.
- `work, *, *, *`: This block is for the network configuration on the job.
- `home, *, *, *`: The configuration for home.
- `*, *, *, *`: This block already exists as well and should be left as is for use in future configurations.

It is easiest to use `YaST` to configure a scheme and then change the name, after running `SuSEconfig`, from `SuSE` to `work` or `home`. Those who wish to edit everything themselves can find more information about terms such as `NETMASK`, `BROADCAST`, etc., and all specific values in PCMCIA-HOWTO (in `/usr/share/doc/packages/pcmcia` you’ll find the current version).

¹Tips can be found in the PCMCIA HOWTO in Section 5.3 “Booting from a PCMCIA device.”

```

# The address format is "scheme,socket,instance,hwaddr" .
case "$ADDRESS" in
SuSE,*,*,*)
    INFO="This scheme is to be configured by YaST/SuSEconfig"
# [... abbreviated ...]
    ;;
work,*,*,*)
    INFO="Network configuration for the company via DHCP"
    IF_PORT=""
    BOOTP="n"
    DHCP="y"
    IPADDR=""
    NETMASK=""
    NETWORK=""
    BROADCAST=""
    GATEWAY=""
    DOMAIN=""
    SEARCH=""
    DNS_1=""
    MOUNTS=""
    start_fn () { return; }
    stop_fn () { return; }
    ;;
home,*,*,*)
    INFO="Network configuration for home, using a fixed address"
    IF_PORT=""
    BOOTP="n"
    DHCP="n"
    IPADDR="10.0.1.23"
    NETMASK="255.255.255.0"
    NETWORK="10.0.1.0"
    BROADCAST="10.0.1.255"
    GATEWAY="10.0.1.1"
    DOMAIN="home.de"
    SEARCH="home.de work.de"
    DNS_1="10.0.1.1"
    MOUNTS=""
    start_fn () { return; }
    stop_fn () { return; }
    ;;
*,*,*,*)
    INFO="Sample private network setup"
# [... abbreviated ...]
    ;;
esac

```

File 8.1.1: /etc/pcmcia/network.opts

```

image = /boot/vmlinuz
root = /dev/hda7
label = work
alias = w
append = xzx SCHEME=work xzx
#
image = /boot/vmlinuz
root = /dev/hda7
label = home
alias = h
append = xzx SCHEME=home xzx
#
image = /boot/vmlinuz
root = /dev/hda7
label = suseconf
alias = s
append = xzx SCHEME=SuSE xzx

```

File 8.1.2: PCMCIA: Example from `lilo.conf`

Scheme changes can be made at boot or run time. The command `cardctl scheme` will tell you which scheme is currently being used. The changes can be made at run time with the same command using the syntax `cardctl scheme <Scheme_name>`. These settings are not lost when restarting the system, which means that the last selected scheme is still active at the next boot unless specifically changed.

This is done either in the ‘boot options’ or directly at the boot prompt (**LILO:**) by adding the following variable to the command: `SCHEME=<Scheme_name>`; more information on giving parameters at the boot prompt can be found in Section 11.3.2 page 239.

You can use the boot manager to create one or more fixed boot schemes. We will continue with our example and create three boot configurations. The `append` parameter in `/etc/lilo.conf` is used for this and other purposes as can be seen in File 8.1.2. After reconfiguring `/etc/lilo.conf` the command `lilo` needs to be called.

Using the example configuration, in the file 8.1.2 by simply adding `w`, `h` or `s` when booting, a specific scheme can be started.

8.1.5 If Things Still Don't Work

Some notebooks have problems with specific PCMCIA cards. Most can be easily solved when you look at the matter systematically. First check to see if the problem lies with the card or the PCMCIA base system. To do this, boot the computer *without* any cards inserted. Only after the base system functions correctly should you insert a card. All important messages can be found in `/var/log/messages`. Therefore this information should be observed with

```
earth: # tail -f /var/log/messages
```

while the necessary tests are running. Doing so can reduce the problem to one of the following two errors.

The PCMCIA base system does not function properly

If the system hangs when booting, after you receive the message "PCMCIA: Starting services:" or other strange things happen, you can prevent PCMCIA from starting at the next boot by entering the parameter `NOPCMCIA=yes` at the boot prompt (`LILLO:`). After the system is running you should load the modules, one after the other, manually, with the commands `modprobe pcmcia_core`, `modprobe i82365` or — in very rare circumstances — `modprobe tcic` and `modprobe ds`. The critical modules are, in each case, the first ones.

If the problem occurs when loading `pcmcia_core` help can be found in manpage for `pcmcia_core` (`man pcmcia_core`). The options described in this can be tested in conjunction with the `modprobe` command.

As an example, we can shut off the APM support of the PCMCIA module; in rare instances, this can lead to difficulties. For this, we can use the option `do_apm`; via `do_apm=0`, the power management will be de-activated:

```
earth:~ # modprobe pcmcia_core do_apm=0
```

If this option is successful, you should set the variable `PCMCIA_CORE_OPTS` in `/etc/rc.config` to:

```
PCMCIA_CORE_OPTS="do_apm=0"
```

Also, from time to time, verifying free IO regions can cause some distress if other hardware components are disturbed by it. You can bypass this via `probe_io=0`. If you want to use a number of options, they should be separated by a space:

```
PCMCIA_CORE_OPTS="do_apm=0 probe_io=0"
```

If the problem occurs when loading `i82365`, help can be found in manpage for `i82365` (`man i82365`).

This error is the result of a resource conflict, which means that an interrupt, an IO port, or a memory area is being shared by two devices. The module `i82365` checks for such sharing but sometimes this check itself leads to problems. Checking the interrupt 12 (PS/2 devices), for example, can lead, on some computers, to the mouse or keyboard being blocked. In this case the parameter `irq_list=<List_of_IRQs>` helps. This list should include all interrupts to be used by the base system and has the following syntax:

```
earth: # modprobe i82365 irq_list=5,7,9,10
```

or permanently in `/etc/rc.config`:

```
PCMCIA_PCIC_OPTS="irq_list=5,7,9,10"
```

In addition to this there are the files `/etc/pcmcia/config` and `/etc/pcmcia/config.opts` which are used by `Cardmanager`. The settings in these files are used for the loading of driver modules for the PCMCIA cards. IRQs, IO ports and memory areas can be assigned or unassigned in the file `/etc/pcmcia/config.opts` as well. The difference between this and the above option is that when resources are unassigned in this file the PCMCIA card does not use them, but the base system module `i82365` still checks them.

The PCMCIA card doesn't function (correctly)

There are three possibilities for errors: the card is not correctly detected, the card causes a resource conflict, or the card is not supported.

If the card is not properly detected, the message "unsupported card in Slot x" will appear in `/var/log/messages`. This message only means that Card-manager cannot correctly assign a driver. `/etc/pcmcia/config` is used for this purpose; it is, so to speak, a "driver database". You can add entries to the "driver database" using existing entries as an example. Using the command `cardctl ident`, you can determine the identification information for the card. More information on this topic can be found in the PCMCIA HOWTO, Section 6 ("Dealing with unsupported cards") as well as the manpage for `pcmcia` (`man pcmcia`). After changing `/etc/pcmcia/config` be sure to reload it with `rcpcmcia reload`.

Generally speaking, it doesn't matter which IRQ or IO port a PCMCIA card uses as long as these are not shared by another device (CD-ROM on 2. IDE controller: IRQ 15; serial or IrDA Port: IRQ 3, IRQ 4; sound, printer: IRQ 5, IRQ 7). In cases of resource conflicts, these should be specifically defined in `/etc/pcmcia/config.opts`. If the problem still exists, these settings can be further specified as a module option in `config.opts`. For example, if the module `pcnet_cs` should use IRQ 5 the following entry would be made:

```
module "pcnet_cs" opts "irq_list=5"
```

Most available options are described in the man pages². If the man page does not include the information needed, there are several options available in manpage for `i82365` (`man i82365`) or the answer lies in the source code of the module.

One common problem with 10/100 Mbit network cards is incorrect detection of the transfer rate. In this case the command `ifport` can be of help. With this the transfer rate can be shown and changed; see also manpage for `ifport` (`man ifport`). The correct setting can be entered in the variable `IFPORT` in `/etc/pcmcia/network.opts`.

8.1.6 Installation via PCMCIA

In some cases PCMCIA support is needed to install SuSE Linux. For this you should choose 'Load PCMCIA Module' from the menu item 'Kernel-Module (Hardware Drivers)' in `linuxrc`. At first two entry fields will appear. Here you can specify options for the modules `pcmcia_core` and `i82365`. Normally these can be left empty. The man pages for `pcmcia_core` and `i82365` can be found as text files on the first CD in the directory `docu`. During installation, system messages are displayed on several virtual consoles and can be viewed using `(Alt) + (F<x>)`.³

After the first part of the installation, the system is partially or totally rebooted. In rare cases the system hangs when starting the PCMCIA system. The installation

²Tip: `rpm -ql pcmcia | grep man` outputs a list all package `pcmcia` man pages.

³In YaST2 you must use `(Ctrl) + (Alt) + (F<x>)` because YaST2 runs under the X Window System.

is far enough along that you can use the boot option **NOPCMCIA=yes** to boot Linux without PCMCIA support, at least in text mode. More information on this can be found in Section 8.1.5 page 188.

8.1.7 Other Help Programs

The program `cardctl` has already been mentioned several times. `cardctl` is an important tool for finding information about PCMCIA devices or taking specific actions. Details can be found in manpage for `cardctl` (`man cardctl`), or simply give the command `cardctl` to receive a list of command options.

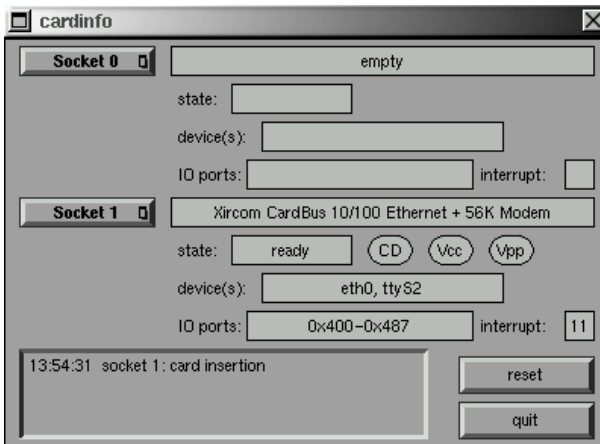


Figure 8.1: PCMCIA — Cardinfo

This program has a graphical front-end, as well, which can be started with the command `cardinfo` (See Figure 8.1). Most options can be controlled with this although, unfortunately, you cannot change between “schemes” using this tool.

Friends of KDE can also use the program `kardinfo`. This tool is basically the same as `cardinfo`.

More help can be found in package `pcmcia` with `ifport`, `ifuser`, `probe` and `rcpcmcia` although they are not often needed in every day problems. To discover all that exists in package `pcmcia`, use the command `rpm -ql pcmcia`.

8.1.8 Compiling PCMCIA

If you update the kernel and do not find a compatible PCMCIA package on the SuSE FTP server, you will then have to compile the PCMCIA package yourself. It is important that the correct kernel is running when you are recompiling, since some information will need to be extracted from there. The package `pcmcia` should already be installed but should not be started; to be safe, run another `rcpcmcia stop`. Then install the PCMCIA source package using YaST and, following that, enter:

```
earth: # rpm -ba /usr/src/packages/SPECS/pcmcia.spec
```

That was it. Now even a new binary packet can be found in `/usr/src/packages/RPMS` which you can use for other systems, given that you implement the same kernel.

8.2 APM — Power Management

8.2.1 Fundamentals

Power management can be used, assuming that the necessary hardware and BIOS routines exist. Most notebooks and desktops include these — the diverse energy saving functions are usually only important for mobile computers. These functions are described below:

Standby – In this operating mode the display is turned off and, by some machines, the processor power is reduced.

Suspend (to memory) – In this mode all system information is written to RAM and the entire system is then suspended. Here the computer uses very little energy; so little, in fact, that the computer can run in this mode from between 12 hours to several days on one battery. The advantage of this is that within seconds you can resume working at the same point at which you left, without rebooting or restarting any programs. This is where using Linux becomes interesting – you never need to turn the computer off – other operating systems become unstable in time. With most modern notebooks you only need to close the lid, suspending the system, and simply open it again later to resume working.

Hibernation (Suspend to disk) – In this operating mode the computer can run right through the winter. That is, all information is saved to the hard disk and the system is turned off. Resuming operation takes about 30–90 seconds and the system is back to the same point where you left off. Some manufacturers offer interesting mixes of suspend and hibernation.

Battery Control – Very interesting.

Automatic Shutdown – Interesting for desktops. After a “shutdown” the computer is completely turned off.

hard disk power off – This function saves not only energy but, for those with a loud hard disk, also stress. Problems can occur, however, with, for example, an editor that automatically saves changes, causing the drive to turn back on, again and again.

Some of these functions are implemented by the BIOS alone. Standby and suspend can be invoked, on many notebooks, with a combination of keys or by simply closing the lid. The operating system has nothing to do with these functions although with the correct kernel and packages installed these can also be invoked with a command. This is helpful if you like to perform certain actions every time the mode is invoked.

8.2.2 The Correct Kernel

For advanced APM functionality a kernel which supports this is needed. Several of the SuSE standard kernel includes this functionality. You can check to see whether these functions are supported with the command `cat /proc/apm`. If a line appears with diverse numbers everything is okay and the command `shutdown -h` should turn the computer off. Because some BIOS's do not correctly support the standards strange things can sometimes happen. Some of these problems can be solved by using a special kernel configurations. More information can be found in our Support Database at <http://sdb.suse.de/sdb/en/html> or on the “Linux Laptop Home Page” at <http://www.cs.utexas.edu/users/kharker/linux-laptop>.

8.2.3 The APM Daemon

The daemon `apmd`, found in package `apmd`, series `ap`, serves to watch the battery and can take specific actions when “standby” or “suspend” is started. This package is not always necessary but can often be of use. In order to start this at boot the variable `START_APM` in `/etc/rc.config` should be set to `yes`. You can also start this with the command `rcapmd start`.

Several variables can be configured in `/etc/rc.config.d/apmd.rc.config`. Because this file contains comments about the settings of these variables we won't go into detail here.

- Although PCMCIA in SuSE Linux is compiled with APM support, you can still sometimes run into difficulties here. Some of the card drivers do not return properly from a suspension (e.g. `xirc2ps_cs`). Therefore, `apmd` can deactivate the PCMCIA system before the suspension and reactivate it afterward. To this end, the variable `APMD_PCMCIA_EJECT_ON_SUSPEND` is set to `yes`.
- If the clock time is no longer correct following a suspension, `APMD_SET_CLOCK_ON_RESUME` is set to `yes`.
- If the X Window System is causing problems, `APMD_LEAVE_X_BEFORE_SUSPEND=yes` helps once the system is up again.
- Likewise the sound modules can be unloaded during suspension. `APMD_STOP_SOUND_BEFORE_SUSPEND=yes` functions for this purpose. Caution: when unloading, sound applications may quit or else the modules were not unloaded. If you are editing audio data, save it first to be safe.

In addition, you can set the “spin-down” behavior of the hard disk differently for the battery mode or even tell the computer to shutdown when the battery reaches a critically low level. Advanced users can add functionalities themselves in the file `/usr/sbin/apmd_proxy`.

8.2.4 More Commands

The package `apmd` has a few other useful programs. With `apm` you can check the actual battery capacity and put the system into standby (`apm -s`) or suspend (`apm -s mode`); see also `manpage` for `apm` (`man apm`).

The command `apmsleep` suspends the system for a certain amount of time; see also `manpage` for `apmsleep` (`man apmsleep`).

To observe a log file without having the hard disk continuously running use the command `tailf` instead of `tail -f`.

There are also tools for the X Window System as well as the command `xapm` which shows a graphic representation of the battery status. If you're using KDE or at least `kpanel` – the program `kbatmon` shows battery status and can suspend or put the system to sleep.

8.2.5 Pause for the Hard Disk

The hard disk can be turned off when not needed with the program `hdparm`. This program also has other useful hard disk functions. The drive can be placed in standby mode with the command option `-y` or completely turned off with `-Y`. And if you enter

```
earth: # hdparm -S 6
```

the drive will be spun down (turned off) after 30 seconds of inactivity. The number in this command is the number of 5 second intervals (in this case 6, so $6 * 5 = 30$). This function is turned off with a value of 0. With larger numbers the interval is also increased, more information can be found in the man pages for this program.

You can set the standby behavior to be dependent on whether the computer is plugged in or in battery mode by configuring `/etc/rc.config.d/apmd.rc.config` but if you do so be sure to set the variable `APMD_CHECK_TIME` to 0.

Because the hard disk is often needed the sleep function is often useless because many programs, for example, text editors, regularly write temporary data to the drive. Programs intended to help, such as `tail -f <logfile>` often cause this as well. Replacing this command with `tailf <logfile>` solves this problem.

8.3 IrDA – Infrared Data Association

IrDA (Infrared Data Association) is an industry standard for wireless communications that uses light in the infrared spectrum. Many laptops delivered today have an IrDA compatible sender/receiver that enables the communication with other devices, such as printers, modems, LANs, or even other laptops. The transfer rate varies from 2400 bps up to 4 Mbps.



Note

Support for this protocol in the Linux kernel is not yet complete. The project is still seen as “experimental” (beta status). Therefore this functionality is not included in the standard kernel. This implies that the driver has not been fully tested yet and may not be as stable as other parts of the kernel.

Software

The package `irda`, series `n` supplies support for infrared ports as well as the IrDA Protocol itself. After installation of this package you can find the documentation under `/usr/share/doc/packages/irda/README`.

From the package `howto` you can install the `IR-HOWTO`. After installation it can be found under `/usr/share/doc/howto/en/IR-HOWTO.gz`. For more information on the Linux IrDA-Project visit the URL <http://irda.sourceforge.net/>.

Configuration

In a normal installation, the variables `START_IRDA`, `IRDA_PORT` and `IRDA_IRQ` are set in the file `/etc/rc.config`. Normally the use of IrDA is deactivated; with the command

```
earth: # rcirda start
```

you can manually activate the port. You can deactivate it with the parameter `stop`. When activated the necessary kernel modules are loaded.

With `IRDA_PORT` (Default `/dev/ttyS1`) and `IRDA_IRQ` (Default `3`) you can configure the IrDA port; they are set by the script `/etc/irda/drivers` when support for the infrared port is activated.

If you give `START_IRDA` a value of `yes` support for the infrared port will be activated at boot. Unfortunately IrDA uses noticeably more electricity due to the so-called “discovery” packets that are broadcast every few seconds, which look for other devices with IrDA communication ports with which to communicate.

Application

In order to print using the infrared port you need to send the print data to the device file `/dev/ir1pt0`. This device file has the same attributes as the normal “wired” port `/dev/lp0`, the only difference being the wireless transmission through infrared light.

You can configure a printer to use this port in the same manner as a parallel port or serial port printer using YaST (see Section 3.7.1 page 70). When printing be sure that the printer is in sight of the computer and that the IrDA support is started.

If you want to communicate with other computers, mobile telephones or the like, address the device using the device file `/dev/ircomm0`. With the Siemens S25 mobile telephone, for instance, you can use the program `wvdial` to access the internet.

Troubleshooting

If devices do not respond when using the infrared port, check to see if the device is found by the computer using the command `irdadump` as `'root'`:

```
earth: # irdadump
```

In the case of a Canon BJC-80 printer in “sight” of the computer `earth`, you’ll receive the following output in regular intervals (see Output 8.3.1).

```
21:41:38.435239 xid:cmd 5b62bed5 > ffffffff S=6 s=0 (14)
21:41:38.525167 xid:cmd 5b62bed5 > ffffffff S=6 s=1 (14)
21:41:38.615159 xid:cmd 5b62bed5 > ffffffff S=6 s=2 (14)
21:41:38.705178 xid:cmd 5b62bed5 > ffffffff S=6 s=3 (14)
21:41:38.795198 xid:cmd 5b62bed5 > ffffffff S=6 s=4 (14)
21:41:38.885163 xid:cmd 5b62bed5 > ffffffff S=6 s=5 (14)
21:41:38.965133 xid:rsp 5b62bed5 < 6cac38dc S=6 s=5 BJC-80 \
hint=8804 [ Printer IrCOMM ] (23)
21:41:38.975176 xid:cmd 5b62bed5 > ffffffff S=6 s=* erde \
hint=0500 [ PnP Computer ] (21)
```

Output 8.3.1: IrDA: `irdadump`

If you receive no output or the device does not respond, check the configuration of the port. Are you using the correct port? Sometimes the infrared port can be found under the device file `/dev/ttyS2` or `/dev/ttyS3` or the device is not being assigned to Interrupt 3. These settings can be changed on nearly every laptop in the BIOS setup.

It is important to note that you can only communicate with devices using IrDA if the device you wish to communicate with supports the protocols `Printer` or `IrCOMM`. With the help of special programs (`irobex_palm3` or `irobex_receive`, please make note of the descriptions in `IR-HOWTO`) you can communicate with devices that use the `IROBEX` Protocol (e. g. 3Com Palm Pilot). Which protocol your device supports can be taken from the output of `irdadump` in brackets after the device name. Support for the `IrLAN` Protocol is still a “Work in progress” and will be included in future versions of Linux.

You can check to see if the IrDA port is actually emitting light with a normal video camera. In contrast to humans, most video cameras can “see” infrared light.

9 Printing

This chapter describes the essential procedures which occur “behind the scenes” when printing takes place.

9.1 Basics and Requirements for Printing

In Linux, data is normally sent to a printer via a “print queue”. The “print spooler” is responsible for queuing the print jobs and sending them to the printer in the correct order, ensuring that they do not interfere with each other.

The data to print often requires conversion to a suitable format before the printer can handle it: graphics files must normally be first converted to a printable format, for example. The “print filter” is a piece of software which is responsible for this conversion.

9.1.1 Different Standard Printer Languages

ASCII text — Most printers can at least print out ASCII text directly. There are two general types of exceptions to this: those printers that first require the text be converted to another of the printer languages listed below and the “GDI printers” which are designed only to work with Windows (TM) and, by default, are capable of printing neither ASCII text or data offered to them in any of the other standard printer languages.

Postscript — Postscript is the standard printer language in Unix and Linux. If no Postscript printer is available, a special filter program (Ghostscript) is used in Linux to convert Postscript data into one of the following standard printer languages.

PCL 5 or PCL 6 — These printer languages are common among laser printers.

PCL 3 — This printer language is used by most HP DeskJet printers and compatible models.

ESC/P, ESC/P2 and ESC/P dot-matrix — These printer languages are used by most Epson printers and compatible models.

9.1.2 The Problem with GDI Printers

Many printers are sold as “Windows printers” or GDI printers. (GDI stands for the Windows Graphical Device Interface: such printers are designed to work

with only one operating system.) They are often difficult or impossible to set up to work with Linux: some of them are capable of using other standard printer languages and are thus usable, while others will only work at all with Windows (TM). Consult the CDB at <http://cdb.suse.de/> or check with the hardware manufacturer if you are unsure.

With GDI printers, the manufacturer does without a standard protocol completely and controls the printer directly with control sequences of the specific model. However, there are printers on the market which can act both as GDI printers and also work with “proper” printer languages.

9.1.3 Hardware and Software Requirements

You should check the following points:

- Check whether your printer is supported by SuSE Linux. Look at the lists on the installed system and on the Internet at:
 - `/usr/share/doc/packages/ghostscript/catalog.devices`
 - <http://cdb.suse.de>
 - <http://www.linuxprinting.org> → ‘The Database’ (<http://www.linuxprinting.org/database.html>)
 - <http://www.cs.wisc.edu/~ghost/printer.html>



Note

You will sometimes hear of printers which require the current version of Aladdin Ghostscript. For licensing reasons, SuSE Linux includes the slightly different GNU Ghostscript. In most cases, a GNU Ghostscript driver can be found which is suitable for any given printer.

- Check that you can make the necessary physical and logical connection to the printer. See Section 9.2 on the next page for details of this.
- A self-compiled kernel could lack the necessary support for printing: if in doubt, first install and boot from a standard SuSE kernel. (This can be done through YaST.)
- It is recommended to install the ‘SuSE Default System’ packages using YaST or the ‘Standard System’, using YaST2. The following packages are necessary along with the ‘SuSE Minimal System’ to print ASCII and Postscript files:
 - package `a2ps`, series `ap`
 - package `aps`, series `ap`
 - package `gs_x11`, series `ap`
 - package `gs_font0`, series `ap`

- package `gs_fonts`, series `ap`
- package `gs_lib`, series `ap`
- package `lprold`, series `n`

For the installation of additional packages, see Section 3.5 page 61.

Possible Problems

- If you have updated an older SuSE Linux, there may be inconsistencies in the printing system’s packages and (configuration) files. In this case, it makes sense to reconfigure the entire printing system. Read the SDB article “Reinstalling a printer” to learn how to proceed in this case (file:///usr/share/doc/sdb/en/html/jsmeix_print-neueinrichten.html).

9.1.4 Defining a Suitable Printer Driver

In the lists mentioned in Section 9.1.3 on the preceding page, Ghostscript drivers for specific printer models are provided. If your supplier cannot give you any Linux-specific information on your printer model, the following tips may help:

- Find out if your printer is compatible with a model which does run in Linux then use the Ghostscript driver for this compatible model.
“Compatible in Linux” means that your printer can print correctly with the same binary control sequences as the compatible model — the printer can “understand” the same printer language directly (rather than just being capable of emulating it in Windows through the use of driver software).
The fact that two printers have similar names is not always a guarantee that they are compatible. This is because printers with similar names do not always understand the same printer language directly.
- The best source of information about your printer’s native language is the manufacturer. In many cases, find the printer language listed among the technical data in the printer manual. Some examples of existing printer languages are given in Section 9.1.1 page 197.
- Ghostscript includes a special `uniprint` driver with parameter files for various printers with the extension `.upp`. If a parameter file exists for your printer, it can be configured in this way.

9.2 Connecting the Printer Locally

9.2.1 Parallel Ports

Connecting a printer to a Linux system is usually done via a parallel port. A printer on a parallel port is addressed via the `parport` subsystem of the kernel.

The parallel ports are made known to the `parport` subsystem by loading additional architecture-specific kernel modules (see Section 7.4.3 page 171). In this way, a number of devices in series can be used *simultaneously* via *one* parallel port. The numbering of device files for parallel port printers starts with `/dev/lp0`. To print via the first parallel port, the modules `parport`, `parport_pc`, and `lp` must be loaded — this is normally performed by `kmod` (*Kernel Module Loader*, see Section 10.2 page 231), as soon as the corresponding device file (e. g., `/dev/lp0`) is accessed.

Configuring the first parallel port

The first parallel port, `/dev/lp0`, is configured by an entry in `/etc/modules.conf` (File 9.2.1).

```
alias parport_lowlevel parport_pc
options parport_pc io=0x378 irq=none,none
# If you have multiple parallel ports, specify them this way:
# options parport_pc io=0x378,0x278 irq=none,none
```

File 9.2.1: `/etc/modules.conf`: First Parallel Port

`io` stands for the first address of the parallel port. The entry `irq=none` puts the parallel port into “polling mode” rather than interrupt driven mode. This is generally a better option and is slower only on very old computers.

For these settings to work, the following settings for the first parallel port must be made in the BIOS or via the computer’s firmware:

- IO address 378 (hexadecimal)
- Interrupt 7 (not relevant in polling mode)
- Mode `Normal` or `SPP` (other modes do not always function)
- DMA is switched off (should be switched to `Normal` mode)

Testing the first parallel port

If the printer is able to print ASCII text, as user ‘`root`’ (system administrator), you should be able to print out a page with the word `hello`, using the command

```
earth: # echo -en "\rhello\r\f" >/dev/lp0
```

Configuring the second parallel port

A second parallel port, which can be addressed via the standard IO address 278 (hexadecimal) (e. g., set by a jumper on an ISA port card), can also be configured with an entry in `/etc/modules.conf` (File 9.2.2 on the facing page).

After a reboot or restarting all services (see page 202), the second parallel port will be available.

```
alias parport_lowlevel parport_pc
# options parport_pc io=0x378 irq=none,none
# If you have multiple parallel ports, specify them this way:
options parport_pc io=0x378,0x278 irq=none,none
```

File 9.2.2: /etc/modules.conf: Two Parallel Ports

Testing the second parallel port

If the printer is able to print ASCII text, as user `'root'`, you should be able to print out a page with the word `hello`, using the command

```
earth: # echo -en "\rhello\r\f" >/dev/lp1
```

9.2.2 Special Slot Cards: ISA-PnP and PCI

If the IO address of an additional parallel port is not known beforehand, this must be determined.

ISA PnP cards

Sometimes fixed values for IO address, interrupt, and mode can be set on such cards, either by jumper or — if you have DOS/Windows — via an included DOS/Windows program. If this is possible, set fixed values for IO address, interrupt, and mode.

Otherwise, the values for IO address, interrupt and mode are only entered on the ISA PnP card when Linux is started. The values entered can either be seen in the Linux boot messages or you can proceed as described in Chapter 7.2 page 166.

PCI cards

The possible IO addresses and interrupts to assign your PCI card are listed with the following command (see Output 9.2.1):

```
earth: # /sbin/lspci -vv | less

...

00:0a.0 Parallel controller: ...
...
Flags: stepping, medium devsel, IRQ 10
I/O ports at b400
I/O ports at b000
I/O ports at a800
I/O ports at a400
...
```

Output 9.2.1: Extract from `lspci -v` for a PCI Port Card

In each case two IO addresses, separated by 400 (hexadecimal), belong together with a parallel port. The configuration entry must appear as shown in File 9.2.3.

```
alias parport_lowlevel parport_pc
# options parport_pc io=0x378 irq=none,none
# If you have multiple parallel ports, specify them this way:
options parport_pc io=0x378,0xb400,0xa800 irq=none,none,none
```

File 9.2.3: /etc/modules.conf: PCI Card with 2 Parallel Ports

After a reboot or restart of all services (as described below), all three parallel ports will become available.

Testing all parallel ports

If you have a normal (ASCII-compatible) printer, you should, as the user 'root', have a page printed out with the word `hello`, using the commands

```
earth: # echo -en "\rhello\r\f" >/dev/lp0
earth: # echo -en "\rhello\r\f" >/dev/lp1
earth: # echo -en "\rhello\r\f" >/dev/lp2
```

Instead of rebooting, it is enough, as the user 'root', to update the list of kernel module dependencies and — if only printers are connected to the parallel ports — stop the printer daemon and unload any kernel modules related to parallel ports:

```
earth: # depmod -a 2>/dev/null
earth: # rclpd stop
earth: # rmmmod lp
earth: # rmmmod parport_pc
earth: # rmmmod parport
```

after this, the kernel modules can be reloaded and the printer daemon restarted:

```
earth: # modprobe parport
earth: # modprobe parport_pc
earth: # modprobe lp
earth: # rclpd start
```

9.2.3 USB Connections

If you intend to use a USB printer, check that USB support is enabled in the BIOS. In an Award BIOS, for example, in the menu 'PNP AND PCI SETUP' (or similar), the entry 'USB IRQ' (or similar) must be set to Enabled (or similar).

Test if the USB printer reacts by entering, as the user root:

```
earth: # echo -en "\rhello\r\f" >/dev/usb/lp0
```

Providing that there is a printer on the first USB port, /dev/usb/lp0, that can print ASCII text, a page with the word `hello` should be printed.

It may sometimes be the case that the USB printer can no longer be addressed, for example, if you remove the USB plug during printing. Then you will need to stop all processes which access /dev/usb/lp0, unload the kernel modules involving

the USB printer, and reload these again by entering the following commands as user root:

```
earth: # fuser -k /dev/usb/lp0
earth: # rcusbmgr stop
earth: # rmmmod printer
earth: # rmmmod usb-uhci      # or:   rmmmod usb-ohci
earth: # rmmmod usbcore
earth: # modprobe usbcore
earth: # mount usbdevfs
earth: # modprobe usb-uhci    # or:   modprobe usb-ohci
earth: # modprobe printer
earth: # rcusbmgr start
```

Check beforehand, with `lsmod`, which USB modules are loaded (either `usb-uhci` or `usb-ohci`) and if further module dependencies exist; for example, the display `usbcore` ... [`printer usb-uhci`]

means that the module `usbcore` is still needed by the modules `printer` and `usb-uhci`. In this case, therefore, the modules `printer` and `usb-uhci` must be removed before the module `usbcore`.

9.2.4 Serial Ports

You can learn about the basics of connecting a printer to a serial port in the “Printing Howto” (`/usr/share/doc/howto/en/Printing-HOWTO.gz`). The following explains some of the details of this process:

1. You should use the `plp` printer spooler (package `plp` from series `n`) instead of the BSD spooler. If you still want to use the BSD spooler, proceed as described further below.
2. Check your `/etc/printcap` file: In addition to the usual entries, there need to be some entries specifying parameters for the serial port:

```
:br#9600:\
:ty=ixon -imaxbel -ixany -ixoff -crtsets:\
```

With `br`, the baud rate of the serial port is specified. `ty` stands for the `stty` options. You may need to adjust both values, depending on the hardware you have. Similarly, you might need to set the correct speed value to match your specific setup — normally one of these values: 2400, 4800, 9600, 19200, 38400, 57600, 115200 and 230400.

With a HP DeskJet 500, for example, the following `stty` settings should do the trick:

```
ixon -ixoff -imaxbel -opost
```

The man page for `stty` explains many of the available options.

3. If you use the BSD spooler: While this printer spooler does read the `br` entry in `/etc/printcap`, it does not know about the `ty` option. However, serial

port parameters can also be configured using the commands **setserial** and **stty**.

To set the parameters to the same values as defined via `/etc/printcap` in the example above, use the following command:

```
earth: # stty 9600 ixon -imaxbel -ixany -ixoff \
-crtscts </dev/ttyS<X>
```

To achieve the desired result, use a sensible value for `ttyS<X>` — this should be `ttyS0` or `ttyS1` in most cases.

Inserting the required **setserial** and **stty** commands into the file `/etc/init.d/boot.local` will ensure that they are run automatically on boot. In addition to that, edit `/etc/printcap`, removing any lines which look like this:

```
:br#9600:\
:ty=ixon -imaxbel -ixany -ixoff -crtscts:\
```

If you did not remove these lines, they would set the wrong default values for the serial port, because the BSD spooler does not know about the **ty** option. Remember that you want the serial port to be configured according to the above commands (possibly at start-up), without them being changed afterwards.

9.3 lpd: the Print Manager and Tools

9.3.1 The Print Manager in the Background

The `lpd` Line Printer Daemon is launched by the `/etc/init.d/lpd` script and runs in the background, assuming the corresponding option **START_LPD=yes** has been set in `/etc/rc.config`.

`lpd` checks `/etc/printcap` to see which print queues are defined. Its job is to organize the execution of spooled jobs:

- it manages local print queues: it sends every job's data file through the appropriate filter — which is assigned by the entry in `/etc/printcap` and by an explicit setting in the control file — and then to the printer port
- it takes care of the order of the jobs
- it checks the `status` of queues and printers in the local spool directories
- it sends jobs to an `lpd` on a remote machine
- it accepts requests from remote machines for local queues or refuses them if authorization fails



Note

Authorization: Only requests from remote hosts listed in `/etc/hosts.lpd` are executed. It is also possible to enter the name in `/etc/hosts.equiv`, but this has *far-reaching* consequences as far as access and security are concerned and, therefore, should be avoided. Additional definitions of the queue can restrict access to certain users or users with an account on the local machine.

9.3.2 Additional Tools and Control Tools for the User

Managing the print queues will often require `'root'` permissions. To remove print jobs of other users, for example, you need to be logged in as `'root'`.

If you are printing a file, for example, via the `raw` queue, which is not available in printer-specific format, the printer will not know what to do with this data: large amounts of paper will be printed with meaningless characters. The same thing will happen if the print filter is not properly configured and produces printer-specific data for an incompatible printer model.

To remove such a print job from the queue, use the following commands to process the print queues.



Note

You need to have `'root'` permissions to manipulate the print queues. Occasionally commands can be run to work with the jobs you sent to the printer without root access.

`lpq` shows your own jobs in the queue. For example,

```
newbie@earth: > lpq -Pqueue
```

```
queue is ready and printing
Rank  Owner   Job  Files      Total Size
active newbie   676  Hello.txt  259420 bytes
1st   newbie   677  letter.dvi  11578 bytes
2nd   newbie   683  picture.gif 37464 bytes
```

Output 9.3.1: Example Output of the Command `lpq`

`lprm` removes your own jobs from a queue. Thus entering this:

```
newbie@earth: > lprm -Pqueue 676
```

will produce the following output:

A print job belongs to the user who initiated it. This user, as well as the job number, are displayed with the `lpq` command. The same job number can appear in different queues, so you always need to give the name of the queue.

```
dfA676Aa05005 dequeued
cfA676Aa05005 dequeued
```

Output 9.3.2: Example Output of the Command `lprm`.

If no job number is given, the currently active job is removed from the relevant queue, if it is your own job.

`lprm -Pqueue` deletes the current job from the specified queue, as long as this belongs to the user who issued the `lprm` command. The printer will still print the data already in its buffer.

`lprm -Pqueue` deletes all print jobs of the user from the specified queue. If `'root'` enters this command, all print jobs are deleted from the specified queue.

Tip

Get more tips and tricks about `lprm` in the “Configuration” manual. Refer to the index of that book.



`lpc` controls the print queues: these can be specified by name, or you can specify `all`. Most of the `lpc` commands listed below can only be executed by `'root'`. The most important options are:

status queue Provides a status report. If the `<queue>` argument is missing, it works like `<all>`, i. e., a status report for all queues is displayed.

disable queue Stops new jobs from being moved to the queue.

enable queue Enables the queue to accept new jobs.

stop queue Stops jobs from the queue being printed (the job currently being printed is completed).

start queue Starts printing jobs from the queue.

down queue Is the same as **disable** plus **stop**.

up queue Is the same as **enable** plus **start**.

abort queue Is the same as **down**, except that a job currently being printed is interrupted immediately. The jobs remain, however, and printing can be continued after the queue is restarted (**up**).

Pass these commands to `lpc` directly in the command line (e. g., `lpc status`) or run `lpc` without parameters so the dialog mode with its own prompt, `lpc>` is started, awaiting the entry of `lpc` commands. End this dialog with `quit` or `exit`.

Note

Print jobs remain in the queues even when you shut down the computer during printing and then start Linux again. A print job which has gone wrong should be removed from the queue with the commands shown above.



9.4 lpr: Sending Print Jobs

The `lpr` is the daily “user interface” for existing print queues. The following procedure takes place:

1. Either the user starts a new print job with `lpr` or an application uses the `lpr` command when the user chooses its ‘Print’ menu item.
2. `lpr` saves the data to the print queue. From there, the print spooler forwards this to the print filter belonging to the print queue.
3. The printer filter converts the data to a format the printer can print directly then sends this to the printer.
4. After the entire print job has been sent to the printer, it is removed from the print queue.

Normally a job is started using:

```
newbie@earth: > lpr text_file
```

or:

```
newbie@earth: > lpr [-Pqueue] text_file
```

If you omit the option `-P` (for “Printer”, the queue’s name), the default is taken from the `PRINTER` environment variable. If this is not set, the default name `lp` is used. The same applies for the commands `lpq`, `lprm`, and `lpc` (see below).

Various print queues are needed for the following reasons:

- If a number of different printers are available, these different printers are addressed via different print queues, such as with


```
newbie@earth: > lpr -Pprinter2 file
```

 for a second printer via the print queue `printer2`.
- Each print queue uses its own print filter. For example, you might use a standard print queue for rapid black and white printing and another print queue (for example, with the name `color`) with a slower print filter which provides higher quality color printing. Access this by


```
newbie@earth: > lpr -Pcolor file
```

If you have configured the printer with `YaST1` or `lprsetup`, the following queues are created by default (to learn about the setup with `YaST2`, refer to the “Configuration” manual):

`lp` — If you print a file via this standard queue, the `apsfilter` is used as the print filter. The filter automatically determines what type of data the file contains and conversion to the printer-specific format is carried out accordingly.

`ascii` — If you print a file via the `ascii` queue, with

```
newbie@earth: > lpr -Pascii file
```

the `apsfilter` is used as the print filter, but this does not automatically define what type of data is contained in the file, rather the command line specifies that the file contains ASCII text and conversion to printer-specific format is carried out accordingly.

`raw` — If you print a file via the `raw` queue, with

```
newbie@earth: > lpr -Praw file
```

then the `apsfilter` is used as the print filter, but this does *not* carry out a conversion to printer-specific format. The file is sent “raw” to the printer. In this case, the file must already contain data in printer-specific format.

In all these cases, the same printer is always used, but it functions differently, depending on the print queue involved.



Tip

For especially large print jobs, it may be useful, instead of making a copy of the data file, to make a symbolic link from the file to the spool directory. This can be done using the `lpr` option `-s`. Remove write permissions to this file until the print job has been completed!

9.5 SETUP (`lprsetup`): Configuring the Print System

The package `aps` installs the configuration program `/var/lib/apsfilter/SETUP`. This is normally started with the command `lprsetup`, symbolically linked to it. It provides the following services:

- A list of all `apsfilter` printer configurations
- Allows insertion and deletion of `apsfilter` printer configurations

9.5.1 Hints on Configuring with `lprsetup`

`lprsetup` is menu-driven and can be operated with the keyboard. Moving between the answer fields ‘OK’ and ‘Cancel’ is done in all menus by pressing `(Tab)`. Use `(↓)` and `(↑)` to move through selection lists.

1. Start `lprsetup` as the user ‘root’ on the text console:

```
earth: # lprsetup
```

Press `(←)` to reach the main menu:

‘EXIT’ *Exit apsfilter setup* — end `lprsetup`.

‘LISTING’ *List all apsfilter entries* — list all existing queues.

‘ENTRY’ *Add/Overwrite/Delete an apsfilter entry* — To add a new printer, the menu item ‘ENTRY’ is used. This leads to another menu, ‘Choose your printer definition’, in which you define your printer.

'DELETE' *Fast delete an apsfilter entry* — To delete an existing apsfilter printer, select **'DELETE'**. You will receive a list of current apsfilter printers to choose from.

2. Select **'ENTRY'**. You will see the following:

RETURN	Back to previous menu
DEVICE	Change printer interface

3. Select **'DEVICE'** (Port).

For printers on a parallel port or a USB, choose **'PARALLEL'**. For ones on a serial port, choose **'SERIAL'**. Parallel ports: `/dev/lp0`, `/dev/lp1`, `/dev/lp2`; USB printers: `/dev/usblp0`, `/dev/usblp1`, `/dev/usblp2`; serial ports: `/dev/ttyS0`, `/dev/ttyS1`, `/dev/ttyS2`.

Confirm your entry with **'OK'**.

4. Select **'PRINTER'** and choose the correct printer type from the following list:

POSTSCRIPT	For a real Postscript printer
HEWLETT-PACKARD	For an HP DeskJet printer
OTHER	Another non-PostScript printer
FREEDDEF	Free definition of driver name

Confirm this with **'COMMIT'**.


Look through the following list for the matching Ghostscript driver or parameter file (recognized by the ending `.upp`) then click on **'OK'**.

5. Enter the resolution. Only leave the defaults if you do not know the correct resolution. If you are unsure, try the resolutions 300, 360, 600, and 720. If you are using a `.upp` parameter file, it is important that the resolution you set here should match that given in the parameter file. Then click **'OK'**.
6. Select **'PAPER'** and go through the following list until you have found the correct paper size. Then click **'OK'**.
7. Select **'COLOR'** and enter **'MONO'** or **'COLOR'** to make it easier for a later configuration to distinguish between black and white and color printing — for this, look at the next item and the following Section 9.6.2 page 212.
8. If you have not yet configured a standard `lp` queue, `lprsetup` will always create the standard queues `lp`, `ascii`, `raw`, and, possibly, `lp-mono`. `lp-mono` is only created if you choose **'COLOR'** in the last item. If the standard queues were already set up, you can, when configuring additional queues, enter your own queue names under **'SPECIAL'** and **'NAME'**. This should only consist of letters and numbers. Short names with lower case letters are preferable, such as `color` for a color print queue.
9. *Important:* With **'ADD'**, the new printer is finally set up. A list is displayed of which print queues were created on the system with the apsfilter configuration. This means there are always several queues created, each enabling a different printer output. Make a note from this list of which print queues were created. Normally only long lines are given. In Output 9.5.1, you should note `color-ascii`, `color`, and `color-raw`. Then choose **'OK'**.

```
### cdjcolor a4dj mono 300 color ###
color-ascii|lp10|cdjcolor-a4dj-ascii-mono-300-color
|cdjcolor a4dj ascii mono 300
color|lp11|cdjcolor-a4dj-auto-mono-300-color
|cdjcolor a4dj auto mono 300
color-raw|lp12|cdjcolor-a4dj-raw-color|cdjcolor a4dj raw
```

Output 9.5.1: List of Print Queues

- Now the configuration files for the `apsfilter` are shown. Any changes should not be made in the global configuration file `/etc/apsfilterrc`, but only in the configuration file allocated to the print queues just created. The latter configuration file has the form `/etc/apsfilterrc.<gs_driver>`, where `<gs_driver>` is the Ghostscript driver chosen in item 4.

If the same Ghostscript driver is chosen for a number of configurations you can still use the same configuration file, since different queues can be distinguished within such a configuration file. For more on this, see Section 9.6.2 page 212. Press  to continue.

- Choose 'RETURN' then 'EXIT'.

9.5.2 Testing the Printer

If the standard queues are created, you should now be able to print the file `/etc/fstab` with the command

```
newbie@earth: > lpr /etc/fstab
```

If the chosen Ghostscript driver allows color printing, test this with:

```
newbie@earth: > lpr /usr/share/ghostscript/*/examples/colorcir.ps
```

9.5.3 Special Cases

With the command `lpr file`, the type of file (e.g., ASCII or Postscript) is automatically defined. If the ASCII file type should be forced, use

```
newbie@earth: > lpr -Pascii file
```

(e.g., to print a Postscript source text as such, as this will otherwise be interpreted as a Postscript file).

For files which already exist in printer-specific format, use:

```
newbie@earth: > lpr -Praw file
```

With

```
newbie@earth: > lpr -Plp-mono file
```

black and white print is forced. To enable this, set the relevant Ghostscript parameters in an `apsfilterrc` file. Refer to Section 9.6.2 page 212.

9.6 The `apsfilter`

9.6.1 How the `apsfilter` Works

When configuring the printer with `YaST1` or `lprsetup`, the following parameters are saved in `/etc/printcap`:

- Ghostscript driver; e. g., `hpdj`, `ljet4`, `stcany.upp`.
- Paper format; e. g., `a4`, `a4dj`, `letter`, `letterdj`.
- Method; `ascii`, `auto`, or `raw`.
- Color; if the method is not `raw`: `mono` or `color`
- Resolution; if the method is not `raw`: e. g., `300`, `360`, `600`, `720`.
- Queue name; if this is explicitly given with `lprsetup` under ‘SPECIAL - Settings for your printer’

The actual print filter, the shell script `/var/lib/apsfilter/apsfilter`, is called up via a link in accordance with the `if` definition in `/etc/printcap`. The above parameters are passed on via this link to the print filter.

```
color-ascii|lp4|hpdj-a4dj-ascii-mono-300-color\
|hpdj a4dj ascii mono 300:\
:lp=/dev/lp0:\
:sd=/var/spool/lpd/hpdj-a4dj-ascii-mono-300-color:\
:lf=/var/spool/lpd/hpdj-a4dj-ascii-mono-300-color/log:\
:af=/var/spool/lpd/hpdj-a4dj-ascii-mono-300-color/acct:\
:if=/var/lib/apsfilter/bin/hpdj-a4dj-ascii-mono-300-color:\
:la@:mx#0:\
:tr=:cl:sh:
```

File 9.6.1: Queue for Color in `/etc/printcap`

Example: If an entry exists in `/etc/printcap` something like what is shown in File 9.6.1, the following parameters are passed on to the `apsfilter`:

```
Ghostscript driver : PRINTER="hpdj"
Paper format      : PAPERSIZE="a4dj"
Method           : METHOD="ascii"
Color            : COLOR="mono"
Resolution       : DPI="300"
Queue name       : QUEUE="color"
```

Depending on how these variables are set, `apsfilter` will behave differently:

METHOD="raw" — The data is sent to the printer device without being converted (e. g., `/dev/lp0`).

METHOD="ascii" — It is assumed that the data is ASCII text and that this text will be converted with `a2ps` to the Postscript format; see the the man page for `a2ps`.

The data in Postscript format is then converted by Ghostscript (`gs`) to the printer-specific format and sent to a printer device. The relevant Ghostscript driver and the resolution are defined in the above variables. Further Ghostscript parameters can be configured with the `apsfilter`.

METHOD="auto" — `apsfilter` automatically tries to define what type of file the data is. The required conversion program is used to transform the data to the Postscript format. The data in Postscript format is then converted by Ghostscript to the printer-specific format and sent to the printer device.

9.6.2 Configuring the `apsfilter`

Through special variables, the behavior of the `apsfilters` can be individually adjusted. These variables are entered in one of the following configuration files:

- `/etc/apsfilterrrc`
- `/etc/apsfilterrrc.<gs_driver>`

where `<gs_driver>` is replaced by the actual name of the Ghostscript driver. These configuration files are loaded by the `apsfilter` in the given sequence — this driver-specific file `/etc/apsfilterrrc.<gs_driver>` overwrites the defaults in the global file `/etc/apsfilterrrc`.

In `/etc/apsfilterrrc`, only global defaults for all print queues should be entered. Special settings should be written in the matching `/etc/apsfilterrrc.<gs_driver>`.

If you configure the printer with YaST2, the parameters are, in part, saved as above in `/etc/printcap`, but especially in `/etc/g.s.upp/y2prn_<queue>.upp`. Here `<queue>` stands for the standard queue `lp` or the name which was assigned with an additional queue (e. g., `color`). The actual print filter `/var/lib/apsfilter/apsfilter` here is also called up via a link in accordance with the `if` entry in `/etc/printcap`. The above parameters are passed on to the print filter partly via this link and partly via `/etc/g.s.upp/y2prn_<queue>.upp`. In this file, the Ghostscript parameters are stored which were set by YaST2 during the printer configuration. Edit this file to enter special Ghostscript parameter values which cannot be adjusted by YaST2. The variables that define the behavior of the `apsfilter` must be entered in one of the following configuration files:

- `/etc/apsfilterrrc`
- `/etc/apsfilterrrc.y2prn_<queue>.upp`

These configuration files are read in by `apsfilter` in above order — the settings in `/etc/apsfilterrrc.y2prn_<queue>.upp` will override the defaults in `/etc/apsfilterrrc`.

In `/etc/apsfilterrc` there are many helpful commentaries on the individual variables. Here is an overview of the most important ones:

FEATURE allows you to define how ASCII texts are printed, currently with the following options:

FEATURE=1 one page of ASCII text per sheet of paper, with headers

FEATURE=2 two pages of ASCII text per sheet of paper, with headers

FEATURE=1n one page of ASCII text per sheet of paper, without headers

FEATURE=2n two pages of ASCII text per sheet of paper, without headers

FEATURE=11 landscape format, with headers

FEATURE=11n landscape format, without headers

GS_FEATURES enables special Ghostscript parameters to be given. The Ghostscript parameters are very varied, depending on the Ghostscript driver used. Information on specific parameter settings can be found in `file:/usr/share/doc/packages/ghostscript/doc/Devices.htm` and `file:/usr/share/doc/packages/ghostscript/doc/hpdj/gs-hpdj.txt`. Refer also to Section 9.8 page 219.

The following continues to refer to the example of the File 9.6.1 page 211.

So that the printer really can always print in black and white in the case of `$COLOR="mono"`, the section from File 9.6.2 should be written into `/etc/apsfilterrc.hpdj`, for example.

```
if [ "$COLOR" = "mono" ]; then
  GS_FEATURES="-sModel=unspec -sColorMode=mono"
else
  GS_FEATURES="-sModel=unspec -sColorMode=CMYK"
fi
```

File 9.6.2: `/etc/apsfilterrc.hpdj`: Settings for Black and White

With the help of the file `file:/usr/share/doc/packages/ghostscript/doc/hpdj/gs-hpdj.txt`, check whether these settings are available for your printer model and, if so, whether they are correct.

Brief information on this:

- For `-sModel=`, choose between 500, 500C, 510, 520, 540, 550C, 560C, 850C, 855C, and `unspec`.
- For `-sColorMode=`, you can choose between `mono`, `CMY`, `CMY+K` and `CMYK`.

Explanations of the various color models (see `file:/usr/share/doc/packages/ghostscript/doc/hpdj/gs-hpdj.txt`):

mono The printer only has a black cartridge, e.g., the models 500, 510 and 520.

CMY The printer can print either with the black cartridge or with the color cartridge. The cartridges usually have to be changed manually, e.g., models 500C and 540.

CMY+K The printer has both a black and a color cartridge, but black and color cannot be mixed. For color printing, black is mixed from the colors, e.g., the models 550C and 560C.

CMYK The printer has both a black and a color cartridge and black and color can be mixed, e.g., the models 850C, 855C, and most other modern deskjets (unspec).

PRELOADS allows a file to be defined which is processed by the `apsfilter` before the actual data is printed. A number of examples of this:

- Adjusting the brightness when printing on color inkjet printers: See `file:/usr/share/doc/sdb/en/html/jsmeix_print-gammakorrektur.html`
- Adjusting page margins when printing Postscript files: See `file:/usr/share/doc/sdb/sdb/en/html/ke_print-margins.html`
- Switching a Postscript printer to duplex printing or choosing a different paper tray with a Postscript printer: See `file:/usr/share/doc/sdb/sdb/en/html/ke_print-duplex.html` and `file:/usr/share/doc/sdb/sdb/en/html/ke_tray.html`

ENCODING allows ASCII text and HTML pages from Netscape to be printed in accordance with ISO-8859-1 or ISO-8859-2 — see also the comments in `/etc/apsfilterrc` and `file:/usr/share/doc/sdb/en/html/jsmeix_print-iso88592.html`.

9.6.3 Finding Errors with `apsfilter`

A protocol recording what passes through the `apsfilter` when a print job is processed can be of help in looking for errors.

Procedure:

1. Log in as user `'root'`.
2. In the file `/var/lib/apsfilter/apsfilter`, remove the comment sign `'#'` in front of `set -x`.
3. From now on, the activities of the print filter are recorded in matching log files which are specified in `/etc/printcap`. For **lf**, the relevant log file of the specific print queues is given, e.g., `/var/spool/lpd/hpdj-a4-auto-color-3` log.

The log file of the last printout in each case is appended to the existing log file, which can lead to these files becoming very large. Empty these log files (without deleting them) using a command such as:

```
earth: # cat /dev/null > /var/spool/lpd/hpdj-a4-ascii-mono-300/log
```

or

```
earth: # cat /dev/null > /var/spool/lpd/hpdj-a4-auto-color-300/log
```

4. After a failed attempt to print, error messages in the log files can help you find the cause of the printing problem. The log files in the above example can be browsed using `less`, for example.

9.7 Printing in the TCP/IP Network

What's it all about?

If you

- use a printer in a TCP/IP network that is attached to another computer
- would like to use a printer via a TCP/IP network that is connected to a “print server unit”
- want to use a printer with a TCP/IP network connection

To accomplish this, you need to filter the print jobs using `apsfilter`.

Preliminary Remarks on Concepts

A computer in a TCP/IP network to which a printer is connected will be called a *print server*, in contrast to a *printserver unit*, which is a small device which connects a printer to a TCP/IP network. A printer that itself has a TCP/IP network connection will be referred to as a *network printer* and the computer that starts the print job as the *client*.

9.7.1 Overview: Specific Cases

The difference is in the filtering of the files to be printed. At some point in the chain, the files need to be converted to a format which the printer can print — in the print language of the printer itself (Postscript, HP PCL, ESC/P). This is done by `apsfilter`. In the case of network printers and printserver units `apsfilter` runs on the client, which sends the information to the network printer or printserver unit. In the case of a print server, `apsfilter` may be run on the client or on the print server itself. There exist, then, several possibilities:

Network printers with filtering at the client:

1. Client converts files to printable files (`-apsfilter->`)
2. Client sends the printable files to the network printer (`-remote->`)

Printserver unit with filtering at the client:

1. Client converts files to printable files (`-apsfilter->`)

2. Client sends the printable files to the printserver unit (`-remote->`)
3. Printserver unit sends the printable files to the printer

A printer connected to a print server with filtering at the client:

1. Client converts files to printable files (`-apsfilter->`)
2. Client sends the printable files to the print server (`-remote->`)
3. Print server sends the printable files to the printer

A printer connected to a print server with filtering at the print server:

1. Client sends the files to the print server (`-remote->`)
 2. Print server converts files to printable files (`-apsfilter->`)
 3. Print server sends the printable files to the printer
- A printer forwarding queue is needed at the place where `-remote->` is. The information is not filtered here, only forwarded.
 - At the place where `-apsfilter->` is, a printer forwarding queue is needed which filters the files. If the filtering is done at the client, this called a “pre-filter” for the “printer forwarding queue”. If the filtering is done at the print server, this is a normal printer forwarding queue such as `lp` or `ascii`.

Requirements:

The printer must be supported by SuSE Linux because the prefilter creates printable files in the same manner as that for a local printer. Refer also to Section 9.2 page 199.

There must be a TCP/IP connection to the network printer and it must be functioning properly.

9.7.2 Configuring a Printer Forwarding Queue on the Client

Log in as system administrator `'root'` on a text console. If you have a graphical login switch to a text console, using `(Ctrl) + (Alt) + (F2)`.

1. Enter the command `/var/lib/apsfilter/SETUP` or `lprsetup`.
2. Press `(←)` to enter the menu `'APSFILTER SETUP'`.
3. Choose `'ENTRY'`, then `'DEVICE'`, then `'REMOTE'`.
4. Enter the IP address of the network printer, the printserver unit, or the print server. The name resolution (DNS) service must be working in both directions — on the client and on the print server side.
5. Enter the description of the forwarding queue on the network printer, the printserver unit, or the print server. In the case of network printer or the printserver units, find the possible descriptions in the documentation of the device. In the case of a print server, it is either `lp` or `ascii`, depending upon which is in use on the print server.

6. Choose 'ADD'.
7. Press \leftarrow .
8. Choose 'RETURN' then 'EXIT'.

Now `remote` is configured on the client. After restarting the printer daemon on the client with the command

```
earth: # rclpd stop
earth: # rclpd start
the queue remote is ready for use.
```

With the command


```
earth: # lpr -Premote <file>
```

the `<file>` is sent without being filtered on the network printer or printserver unit. If the printer is connected to a print server, the file `/etc/hosts.lpd` must include the IP address of the client. If name resolution functions, you may also enter the name of the client here.

After restarting the printer daemon on the print server it will accept print jobs from the given client. If the printer is connected to the print server and the filtering is done on the print server, you are now done with the configuration.

9.7.3 Configuring a Prefilter on a Client

1. Start `lprsetup`.
2. Press \leftarrow .
3. First choose 'ENTRY', then 'DEVICE', then 'PREFILTER'.
4. Choose the configured queue `remote`.
5. Choose 'PRINTER'.
6. Select the relevant print format from the list 'POSTSCRIPT', 'HEWLETT-PACKARD', 'OTHER', or 'FREEDDF'.
7. Choose 'COMMIT'.
8. Look through the following list until you find the correct Ghostscript driver or parameter file then choose 'OK'.
9. Enter the correct resolution then choose 'OK'.
10. First press 'RETURN' then 'PAPER'.
11. Enter the correct paper size then choose 'OK'.
12. Choose 'COLOR'.
13. Choose either 'MONO' or 'COLOR' and then press 'OK'.
14. Choose 'ADD'. You will receive output telling you which queue will be used on your system with the `apsfilter` configuration. Make a note of this then choose 'OK'.

15. Press .
16. Press first 'RETURN' then 'EXIT'.

After restarting the printer daemon on the client with the commands

```
earth: # rclpd stop
earth: # rclpd start
```

your new prefilter queue on the client is ready for use.

With the command

```
earth: # lpr -Pprefilter-queue <file>
```

the <file> is sent to the `prefilter-queue` given. Once there, it is converted by `apsfilter` into printable files and, through the `remote` queue, is sent to the network printer, printserver unit, or print server.

9.7.4 Problems with Configuration

Checking the prefilter configuration

Connect the printer to the first parallel port on the computer to use and configure the printer, just for test purposes, locally on the computer.

To do this, proceed with Section 9.2 page 199.

This tests whether the problems lie on the printer or the network. If the printer functions, you know the Ghostscript driver and the other parameters for the configuration were correct. Now, configure your printer as a network printer, in accordance with the above instructions.

The network printer or printserver unit does not work properly

There are occasionally problems with the print spooler, for example, when more than one printer is attached or when more than one file at a time needs to be printed. Since these are connected with the print spooler in the printserver unit or in the network printer, nothing much can be done. It is therefore generally easier to connect the printer to a print server.

The following remedy may help, however: a single computer serves as a print server for the printserver unit or the network printer. All clients send print jobs to this print server and only this print server takes on the spooling of print jobs and sends these on to the printserver unit or the network printer. Filtering can take place on the client or on the print server. Details of this can be found at: file:///usr/share/doc/sdb/en/html/jsmeix_print-netzwerk.html.

Background information:

The “printer forwarding queue” and the “prefilter queue” cannot be contained in the same queue because the print daemon `lpd` (from package `lprold`) ignores distant queue filter information given in `/etc/printcap`. A prefilter needs its own queue.

9.8 Some Words on Ghostscript

If you do not own a Postscript printer, Ghostscript is the most popular choice for a filter. Ghostscript accepts Postscript files and contains many printer drivers for conversion purposes. Ghostscript is much more, however, than just a collection of useful printer drivers — Ghostscript can also prepare Postscript files for screen output (for gv) or convert them to PDF format.

Ghostscript is a very comprehensive tool featuring many command line options. Find a detailed description of this program in `file:/usr/share/doc/packages/ghostscript/doc/Use.htm`. If invoked directly, Ghostscript presents its own input prompt `GS>` after your command has been processed. Exit this prompt by entering `quit`.

The help command

```
newbie@earth: > gs -h | less
```

lists the most important options, the version number of the program, and the *current list of supported devices*. The `uniprint` details also appear here — the parameter files for `uniprint` are listed in `file:/usr/share/doc/packages/ghostscript/catalog.devices`.

9.8.1 Examples of Working with Ghostscript

A number of Postscript files can be found in the directory `/usr/share/doc/packages/ghostscript/examples`.

Among those, `file:/usr/share/doc/packages/ghostscript/examples/colorcir.ps` is one of the files suitable for a printer test. In X, display a Postscript file on the screen with the command `gs`:

```
newbie@earth: > gs /usr/share/doc/packages/ghostscript/examples/colorcir.ps
```

To close this, press `(Ctrl) + (C)` in the terminal window from which you started `gs`.

Converting a Postscript file to the printer-specific format for a PCL-5 or PCL-6 printer can be performed, for example, with the command

```
newbie@earth: > gs -q -dNOPAUSE -dSAFER -sOutputFile=/tmp/out.prn \
-sDEVICE=ljet4 -r300x300 \
/usr/share/doc/packages/ghostscript/colorcir.ps quit.ps
```

where the command should be written in a single line, and the backslash (`'\'`) suppressed. Furthermore, it is assumed that the file `/tmp/out.prn` does not yet exist.

Converting a Postscript file to the printer-specific format for an HP DeskJet (color printer with PCL 3) can be done with one of the following commands:

```
newbie@earth: > gs -q -dNOPAUSE -dSAFER -sOutputFile=/tmp/out.prn \
-sDEVICE=hpdj -r300x300 \
-sModel=500 -sColorMode=mono -dCompressionMethod=0 \
/usr/share/doc/packages/ghostscript/examples/colorcir.ps quit.ps
newbie@earth: > gs -q -dNOPAUSE -dSAFER -sOutputFile=/tmp/out.prn \
-sDEVICE=cdjmono -r300x300 \
/usr/share/doc/packages/ghostscript/examples/colorcir.ps quit.ps
```

```

newbie@earth: > gs -q -dNOPAUSE -dSAFER -sOutputFile=/tmp/out.prn \
-sDEVICE=cdj500 -r300x300 \
/usr/share/doc/packages/ghostscript/examples/colorcir.ps quit.ps
newbie@earth: > gs -q -dNOPAUSE -dSAFER -sOutputFile=/tmp/out.prn \
-sDEVICE=cdj550 -r300x300 \
/usr/share/doc/packages/ghostscript/examples/colorcir.ps quit.ps
(Every command is to be entered on each line without '\'.)

```

The conversion of a Postscript file to the printer-specific format for an ESC/P2 or ESC/P or ESC/P dot matrix printer takes place using one of the following commands:

```

newbie@earth: > gs -q -dNOPAUSE -dSAFER -sOutputFile=/tmp/out.prn \
@stcany.upp \
/usr/share/doc/packages/ghostscript/examples/colorcir.ps quit.ps
newbie@earth: > gs -q -dNOPAUSE -dSAFER -sOutputFile=/tmp/out.prn \
-sDEVICE=stcolor -r360x360 \
-dBitsPerPixel=1 -sDithering=gsmono -dnoWeave -
sOutputCode=plain \
/usr/share/doc/packages/ghostscript/examples/colorcir.ps quit.ps

```

Here you can see the difference in the command when using a parameter file for the uniprint driver and when using one of the other Ghostscript drivers. Since all driver-specific parameters are given in the uniprint parameter file, no further driver-specific parameters need to be given, in contrast to the other Ghostscript drivers.

After each of the above commands, the printer-specific data is located in `/tmp/out.prn`, which can be sent by `'root'` directly to the printer with the following command

```
newbie@earth: > cat /tmp/out.prn >/dev/lp0
```

assuming that the printer is connected to the first parallel port, `/dev/lp0`.

If an ASCII text file is printed with Ghostscript, this should first be converted to the Ghostscript format with

```
newbie@earth: > a2ps -1 --medium=A4dj --output=/tmp/out.ps textfile
```

and then, with

```
newbie@earth: > gs -q -dNOPAUSE -dSAFER -sOutputFile=/tmp/out.prn \
<driver--specific Ghostscript parameter> \
/tmp/out.ps quit.ps
```

converted to the printer-specific format, which is then, as `'root'`, sent to the printer with

```
newbie@earth: > cat /tmp/out.prn > /dev/lp0
```

Find all the latest information on the *Ghostscript Homepage* (versions, printer support, etc.). This is located at <http://www.cs.wisc.edu/~ghost/>.

9.9 Examples of Your Own Print Filters

9.9.1 Preliminary Notes

The aim of this section is not to present an alternative to the `apsfilter`, but rather to explain the background to printing in Linux, using the example of a simple

self-created print filter. This means that, in special cases, it is also possible to use an additional print queue with your own printer filter if the `apsfilter` cannot be properly configured for this special case. The examples given here are kept as simple as possible to illustrate the essential steps. This is why things like error correction measures in the bash scripts were left out.

We are assuming here that the printer is connected to the first parallel port, `/dev/lp0`. If the printer is connected to a different port on the computer, then you should use the device of this port, rather than `/dev/lp0` (see Section 9.2 page 199).

A print filter receives the data to print from the print spooler, via the standard input. The print filter must convert this data into the printer-specific format then output this via the standard output. The print spooler ensures that everything the print filter sends to the standard output arrives at the printer device, `/dev/lp0`. The kernel in turn forwards everything which arrives at the printer device onto the defined port (e.g., to the IO address `0x378`). The hardware ensures that everything that is sent to the I/O address `0x378` is sent via the parallel connection to the printer. The printer interprets this data stream and prints accordingly.

The following commands can usually only be carried out as the user `'root'`, because normal users may not directly access the printer device unless the user `'root'`, with the command

```
earth: # chmod a=rw /dev/lp0
```

has given all users permission to access the printer device directly.

Commands can be given as follows:

```
earth: # cat ASCII-Textfile >/dev/lp0
```

Here it should be clear that `ASCII-text_file` needs to be replaced by the name of an existing ASCII text file.

9.9.2 A Simple Example of the Basic Method of Working

Through the command

```
earth: # echo -en "\rhello\r\f" >/dev/lp0
```

no print spooler or print filter becomes active, since the printer device `/dev/lp0` is addressed directly. Through this, only the ASCII characters `'\r'`, `'h'`, `'e'`, `'l'`, `'l'`, `'o'`, `'\r'`, and `'\f'` are sent directly to the printer. The characters `'\r'` and `'\f'` are carriage return and formfeed characters in ASCII notation, causing the printer to perform the corresponding actions. As long as the printer can print ASCII characters directly, it will print the word `hello` and issue the page.

With

```
earth: # cat ASCII-Textfile >/dev/lp0
```

no print spooler or print filter will become active, since the printer device `/dev/lp0` is directly addressed. The ASCII characters from the ASCII text file are sent directly to the printer.

In Linux, two ASCII text lines are separated only by a linefeed ASCII character (line break). In the various DOS versions and in Windows (hereafter referred

to as DOS/Windows) two ASCII text lines are separated by a line feed ASCII character and a carriage return ASCII character.

If you send an ASCII test file directly to the printer with the above command, this will normally appear on the page as follows:

```
First Line
      Second Line
                Third Line
                        ...
```

because the printer will perform a line return, but not a carriage return, since no ASCII carriage return character was sent.

You can, however, adjust the printer so that it carries out both a line feed and a carriage return when it receives an ASCII line feed character.

For printers which understand the PCL 3 printing language (HP DeskJets and compatible models), the printer can be adjusted with the escape sequence `\033&k2G`, so that it performs a line feed and a carriage return when receiving an ASCII line feed character.

With

```
earth: # echo -en "\033&k2G" >/dev/lp0
```

the escape sequence is sent to the printer then with

```
earth: # cat ASCII-Textfile >/dev/lp0
```

the ASCII text file is printed with the correct line wrapping. You may still need to enter

```
earth: # echo -en "\f" >/dev/lp0
```

to have the last page removed from the printer.

Usually, umlauts will not be correctly printed, however, since they are coded differently in DOS/Windows than in Linux and printers are normally set by default for DOS/Windows.

With

```
earth: # cp ASCII-Textfile ASCII-Textfile.ibmpc
```

```
earth: # recode latl..ibmpc ASCII-Textfile.ibmpc
```

the ASCII text file is first copied to `ASCII-text_file.ibmpc` then `ASCII-text_file.ibmpc` is newly coded to conform with DOS/Windows.

If you now enter

```
earth: # cat ASCII-Textfile.ibmpc >/dev/lp0
```

both line wraps and umlauts should be printed correctly. Because both line wraps and umlauts in the `ASCII-text_file.ibmpc` are now coded in accordance with DOS/Windows, a special escape sequence is no longer needed to adjust the appropriate line wrapping behavior for the printer.

So, with

```
earth: # cp ASCII-Textfile ASCII-Textfile.ibmpc
```

```
earth: # recode latl..ibmpc ASCII-Textfile.ibmpc
```

```
earth: # cat ASCII-Textfile.ibmpc >/dev/lp0
```

any ASCII text file can be correctly printed on any printer which can print ASCII text.

If this works, the next step is to create a print filter which automatically takes care of this conversion of the ASCII text file into the printer-specific format.

Following the pattern of `/var/lib/apsfilter`, a subdirectory for your own print filter is created and you change to this subdirectory (as the user 'root'):

```
earth: # mkdir /var/lib/myprinterfilter
earth: # cd /var/lib/myprinterfilter
```

Create a bash script (as an ASCII text file) with the name `asciifilter`, as shown in File 9.9.1.

```
#!/bin/bash

# make a temporary file
INPUT="$(mktemp /tmp/asciifilter.$$XXXXXX)"

# First store everything from stdin in $INPUT
# to have the input as a regular file
cat > $INPUT

# Recode the INPUT
recode lat1..ibmpc $INPUT

# Add a FormFeed at the end of $INPUT
# to get the last page out of the printer
echo -en "\f" >> $INPUT

# Send $INPUT to stdout
cat $INPUT

# Remove the INPUT file
rm $INPUT
```

File 9.9.1: `/var/lib/myprinterfilter/asciifilter`

Make this bash script readable and executable for all users, with

```
earth: # chmod a+rx /var/lib/myprinterfilter/asciifilter
```

Create a new print queue with `lprsetup`. Select any Ghostscript driver which you would normally not use. Resolution and paper size are not relevant. Select a driver which you would not normally use 'MONO'. But for 'SPECIAL Settings for your printer' enter, under 'NAME Optional name for the queue', `af` as the name, which should stand for `ascii filter`, assuming you do not already have a print queue with this name. The queues are then created by `lprsetup: af-ascii, af, and af-raw`.

Below, only the `af` queue is used. The queues `af-ascii` and `af-raw` are actually irrelevant and probably cannot be used, since the Ghostscript driver selected will not function for the printer connected.

In `/etc/printcap` you will now find an entry similar to the one shown in File 9.9.2 on the next page (abbreviated).

Stop the print spooler with

```
af...
      :if=/var/lib/apsfilter/bin/...-af:\
      :la@:mx#0:\
      :tr=:cl:sh:
```

File 9.9.2: /etc/printcap: Your Own Filter

```
earth: # rclpd stop
```

and just change the line in the above entry in /etc/printcap

```
      :if=/var/lib/apsfilter/bin/...-af:\
```

to

```
      :if=/var/lib/myprinterfilter/asciifilter:\
```

then, after saving /etc/printcap, start the print spooler again with

```
earth: # /sbin/init.d/lpd start
```

and now, with

```
earth: # lpr -Paf ASCII-Textfile
```

every user should be able to print via the new af queue.

This new queue does exactly the same as the apsfilter, as long as you set `USE_RECODE_NOT_A2PS="yes"` in an `apsfilterrc`.

To make sure that data is not sent unintentionally to the printer via the irrelevant queues, `af-ascii` and `af-raw`, the lines for these queues should, in each case, be changed from

```
      :lp=/dev/lp0:\
```

to

```
      :lp=/dev/null:\
```

9.9.3 A More Complex Example

Concern

The existing driver is not a Ghostscript driver. In this case you eventually have to create your own printer driver. You have to distinguish between two variations:

1. There is no Ghostscript driver available for this printer (e. g., the printer is a GDI printer).
2. A Ghostscript driver is available for this printer, but there is another non-Ghostscript driver available which creates better output (e. g., a driver specially optimized for this printer).

Such drivers are normally used as an addition to the Ghostscript driver by converting the Ghostscript output to the printer specific format. As this is the most complex situation, this example can be used to create a separate printer driver later on.

Requirements

It is assumed that a driver exists with the current version of SuSE Linux or that a driver can be downloaded from the Internet. It is also assumed that you can handle Unix source (e.g., .zip or .tar.gz archives and .rpm packages); see also Section 12.3 page 278. After unpacking such an archive, you normally find installation instructions in files like README or INSTALL or in a subdirectory named doc.

Background Information

In the case of a .tar.gz archive, it is normally necessary to compile and install the driver. We assume that the driver is being installed as /usr/local/bin/printerdriver here.

Normally, the driver simply converts the Ghostscript output into a printer-specific format (see Section 9.8 page 219). Therefore, Ghostscript needs to be called like this to create a 600 dpi resolution image:

```
-sDEVICE=pbmraw -r600x600
-sDEVICE=pngmraw -r600x600
```

Have a look at the documentation provided with the driver to find out exactly which Ghostscript driver and which resolution is needed.

WATCHME-O.N. a translation of this paragraph was not included in the old e version. we translate it but leave it commented out... For the purposes of print filtering, Ghostscript does not take its input directly from the file to be printed, but from standard input. To make this work, a single - must be included in the command. Similarly, to enable Ghostscript to send its output to the filter program proper, include the parameter **-sOutputFile=-**.

The output of Ghostscript is forwarded to the driver via a pipe:

```
| /usr/local/bin/printerdriver Parameter
```

Check the documentation of the driver to find out which parameters are necessary.

The complete command looks like this

```
gs -q -dNOPAUSE -dSAFER -sOutputFile=- \
  -sDEVICE=<device> -r<resolution> \
  - quit.ps \
  | /usr/local/bin/printerdriver <parameter>
```

The variables <device>, <resolution>, /usr/local/bin/printerdriver, and <parameter> need to be replaced by the values described in the driver documentation.

For a printer supported by Ghostscript, the command line is simplified to

```
gs -q -dNOPAUSE -dSAFER -sOutputFile=- \
    -sDEVICE=<ghostscript-driver> -r<resolution> \
    <opt-parameter> \
    - quit.ps
```

where <ghostscript-driver>, <resolution>, and <opt-parameter> need to be replaced by valid values (refer to Section 9.8 page 219).

As Ghostscript expects input in the Postscript format, the data needs to be converted to Postscript before Ghostscript is called.

Normally, it will only do a conversion from ASCII text to Postscript, as most applications create either Postscript (e.g., Netscape and StarOffice) or ASCII text (e.g., editors). The format of the data can be identified by the **file** command.

ASCII text can be converted to Postscript with the **a2ps** command. To get details on the various options of this tool, refer to the man page for **a2ps**.

Installation and usage

Create a directory `/var/lib/myprinterfilter` and change to that directory. Create a bash script called `printerfilter` (refer to 9.9.3 on the next page).

You can use this example as it is. All you need to change is the variable **GSDEVICE** and **DRIVER**. Change them to fit to your printer. For a printer supported by Ghostscript enter:

```
DRIVER='cat -'
```

For a Postscript printer, that requires no Ghostscript, enter:

```
GS='cat -'
DRIVER='cat -'
```

Change the variables **A2PSLAYOUT** and **A2PSENCODING** as desired. For additional information, refer to the man page for **a2ps**.

Make this bash script executable to everyone on the system and create a printer queue with **lprsetup**. Have a look at Section 9.9.2 page 221 for details on how to do that.

If the printer filter does not work as expected, have a look at the log file which is created during a print job and records all the output of the commands. Refer to Section 9.6.3 page 214 to read more about this.

```

#!/bin/bash

# For debugging see for stderr in /var/spool/.../log
#set -x

# Define the program calls for the filter pipes

# The a2ps call
A2PSLAYOUT="-1 --center-title= --borders=yes"
A2PSENCODING='--encoding=ISO-8859-1'
A2PS="a2ps $A2PSLAYOUT $A2PSENCODING --output=-"

# The gs call
# in GSDEVICE Ghostscript device and parameters have to be
# set
GSDEVICE='-sDEVICE=Device -rResolution'
GS="gs -q -dNOPAUSE -dSAFER -sOutputFile=- $GSDEVICE - quit.ps"

# The call for an optional printerdriver
DRIVER='/usr/local/bin/printerdriver Parameter'

# make a temporary file
INPUT="$(mktemp /tmp/printerfilter.$$XXXXXX)"

# First store everything from stdin in $INPUT
# to have the input as a regular file
cat > $INPUT

# Determine the FILETYPE of the data in $INPUT
FILESTRING=$(file $INPUT | tr 'A-Z' 'a-z')
case "$FILESTRING" in
  *postscript*) FILETYPE=postscript ;;
  *ascii*|*text*|*english*|*script*) FILETYPE=ascii ;;
  *) FILETYPE=unknown ;;
esac

# Do the filter pipe according to the FILETYPE
case "$FILETYPE" in
  postscript)
    cat $INPUT | eval $GS | eval $DRIVER ;;
  ascii)
    cat $INPUT | eval $A2PS | eval $GS | eval $DRIVER ;;
  unknown)
    echo "Unknown filetype: $FILESTRING" \
      | mailx -s "printerfilter" root ;;
esac

# Remove the INPUT file
rm $INPUT

```

File 9.9.3: A Complex Printer Filter

10 The Kernel

The standard SuSE kernel, which is written to disk after installation (and which is found in a correctly installed system under `/boot`), is configured to support as wide a range of hardware and other kernel features as possible. For this reason, this kernel is not specifically tuned to your own hardware. If you have made the right selection during the installation or update, then this kernel will be specifically optimized for your processor. However, it is possible that this kernel may not be ideal for your purposes, most of which relate to network server and routing tasks. In addition, a small amount of RAM will be wasted as a small number of unnecessary drivers are loaded.

Several `Makefiles` are provided with the kernel to help automate the process. These `Makefiles` handle nearly all the details for you. The only thing you have to do by hand is to make selections which correspond to hardware settings and other kernel features.

Installation support is not provided for kernels that are not included with SuSE Linux; we will be pleased to help you, however, in the context of our Professional Services.



Note

The description below is based on kernel series 2.4x and 2.2.x. Many of the issues mentioned here will be valid for 2.0.x, but they might differ in the details!

10.1 Kernel Sources

To build a kernel, the following packages must be installed: the kernel sources, (package `kernel-source`), the C compiler (package `gcc`), the GNU binutils (package `binutils`) and the include files for the C compiler (package `glibc-devel`). They are located in series `D` (Development). It is highly recommended to install the C compiler in any case, since the C language is inseparable from UNIX operating systems.

The kernel sources are located in `/usr/src/linux`. If you plan to experiment with different kernel sources, you can unpack them in different directories under the directory `/usr/src` and create a symbolic link, `/usr/src/linux`, to the current kernel source. This is what YaST does automatically.

Since there is a lot of software that relies on the sources being in `/usr/src/linux`, you should maintain this directory as a symbolic link to your current

kernel source to provide an error-free compilation of system programs which need to access the kernel sources.

10.2 Kernel Modules

Many drivers and features no longer have to be compiled directly into the kernel, but can be loaded at runtime via kernel modules. Which drivers are to be compiled into the kernel and which are loaded as runtime modules is defined in the kernel configuration.

Kernel modules are located at `/lib/modules/<version>`, where `<version>` is the actual kernel version.

The use of this feature decreases the size and RAM requirements of the kernel, which is quite desirable. For this reason, it is best to use this feature wherever it is reasonable. Basically, components which are not required for the system to boot up may be built as modules. This makes sure that the kernel doesn't get too big to be loadable by the BIOS or a bootloader. Drivers which are required for boot, such as `ext2`, the SCSI drivers on a SCSI-based system, and similar items should be compiled into the kernel. In contrast, items such as `isofs`, `msdos`, or `sound`, which are not needed for starting up your computer system, may be built as modules.

Handling Modules

The following commands are available for your use:

- **insmod**
insmod loads the requested module after searching for it in a subdirectory of `/lib/modules/<version>`. It is preferable, however, to use **modprobe** (see below) rather than **insmod**, which should no longer be necessary.
- **rmmod**
Unloads the requested module. This is only possible if this module is no longer needed. It is not possible to unload the `isofs` module (the CD-ROM file system), for example, while a CD is still mounted.
- **depmod**
Creates the file `modules.dep` in `/lib/modules/<version>` where dependencies of all of the modules are defined. This is necessary to ensure that all dependent modules are loaded with the selected ones. If **START_KERNELD** is set in `/etc/rc.config`, this file is created each time the system is started.
- **modprobe**
Loads or unloads a given module while taking into account dependencies of this module. This command is extremely powerful and can be used for a lot of things (e. g., probing all modules of a given type until one is successfully loaded). In contrast to **insmod**, **modprobe** checks `/etc/modules.conf` and is the preferred way for loading modules. For detailed information on this topic, please refer to the corresponding manual page.

- **lsmod**

Shows you which modules are currently loaded and by how many other modules they are being used. Modules started by the kernel daemon are tagged by `autoclean`, which shows that these modules will be removed automatically when they reach their idle time limit.

`/etc/modules.conf`

Loading of modules is further influenced by `/etc/modules.conf`. See also `manpage` for `depmod` (`man depmod`).


In this file, the parameters for modules which access hardware directly can be entered as such modules may need system-specific options (e. g. CD-ROM driver or network driver). The parameters entered here are, in principle, identical to those given to the boot prompt of the kernel e. g. for LILO (see Section 11.3.2 page 239), but in many cases the names which are used at the boot prompt are different (to compare these, see Section 11.3.4 page 251). If a module failed to load, try specifying the hardware in this file and use `modprobe` to load the module instead of `insmod`.

Kmod — The “Kernel Module Loader”

>From version 2.2.x, the kernel module loader is the most elegant way to use modules and replaces the old kernel daemon (`kerneld`). This kernel feature allows the kernel to launch `modprobe` directly and ensures that the necessary modules are loaded as soon as required by the kernel.

To use the kernel module loader, you must set the corresponding variable in the kernel configuration ‘Kernel module loader’ (`CONFIG_KMOD`)

The drivers needed to access the root file system should be compiled directly into the kernel. So you should not configure your SCSI driver or your file system (normally: `ext2`) as modules!



Caution

Because SuSE Linux however, now uses `initrd` (initial ramdisk) and integrates the SCSI driver, for example, using this method, you must make sure when compiling your own kernel that you adjust the variable `INITRD_MODULES` in the file `/etc/rc.config` (Section 14.6 page 315) and comment out the `initrd` line in `/etc/lilo.conf` (see Section 13.4.5 page 291). If you don't do this the kernel will hang when you boot.

Kernel modules come in handy for rarely used functionality, such as `parport` and printer support, drivers for floppy drives and file systems that are rarely used.

Kmod is not designed to automatically unload modules; the potential saving in memory is only marginal for the RAM capacity of computers today; see also

`/usr/src/linux/Documentation/kmod.txt`. For reasons of performance, it is better for server machines which have special tasks to perform and need only a few drivers to have a “monolithic” kernel.

10.3 Kernel Configuration

The configuration of the kernel that was set up during installation or during an update can be taken from the file `/usr/src/linux/.config` (see Section 2.1.8 page 18).

Configuring the kernel can be done in three different ways:

1. On the command line
2. In a menu in text mode
3. In a menu in the X Window System

Here is a short overview of these three methods.

Tip

If you want to use an already existing `.config`, just type `make oldconfig` and then proceed as described in Section 10.5 on the next page.



Configuring on the Command Line

To configure the kernel, just change to `/usr/src/linux` and enter:

```
earth:/usr/src/linux # make config
```

You are asked to choose the options that you want supported by the kernel you are about to build. There are two or three possible answers here: `(y)`, `(n)` or `(m)`. ‘m’ means that this device is not compiled directly into the kernel but as a module instead. Any driver that is needed to boot the system should be integrated into the kernel, and not be loaded as a module. If you press any other key, you get a short help text about the current option.

Configuring in Text Mode

A much more convenient way of configuring the kernel can be achieved by typing:

```
earth:/usr/src/linux # make menuconfig
```

With `make menuconfig`, you can review your changes, go through the questions in your own preferred order and, in the event of a mistake, you do not have to go through all the questions again.

Configuring in the X Window System

If you have installed and configured the X Window System (package `xf86`) and Tcl/Tk (package `tcl` and package `tk`), you can use:

```
earth:/usr/src/linux # make xconfig
```

as an alternative. You will be presented with a GUI (Graphical User Interface) which makes kernel configuration very user-friendly. You should have started the X Window System as `'root'` or you will have to take additional steps into consideration (e. g., taking over the display from another user).

10.4 Settings in the Kernel Configuration

All the individual configuration possibilities of the kernel cannot be covered here in detail. Please make use of the numerous help texts available on kernel configuration. The latest kernel documentation is always in `/usr/src/linux/Documentation`.

10.5 Compiling the Kernel



Tip

You may remove the comment in the main `Makefile` (app. line 92). That's the line containing: `export INSTALL_PATH=/boot`. This lets you install your own kernel to `/boot`.

We recommend that you compile a “bzImage”. As a rule this avoids the problem of the kernel getting *too large*, as can easily happen if you select too many features and create a “zImage” (messages such as “kernel too big” or “System is too big”) are then typical.

After adapting the kernel configuration to your needs, start compilation by entering:

```
earth:/usr/src/linux # make dep
earth:/usr/src/linux # make clean
earth:/usr/src/linux # make bzImage
```

These three commands can be entered on one line as well. They are started one after the other. This might be useful if you want to compile a kernel overnight, for example. Just enter:

```
earth:/usr/src/linux # make dep clean bzImage
```

Depending on your system, it now takes just a few minutes (AMD Athlon / Intel Pentium III) up to several hours for a 386 with 8 MB.

While compiling, you can still work on one of the other consoles.

After a successful compilation, you will find the compressed kernel in `/usr/src/linux/arch/i386/boot`. The kernel image – the file which contains the

kernel — is called `bzImage`. If you cannot find this file, in all probability an error has occurred during the kernel compilation. This can easily be overlooked due to the large amount of screen output. If an error has taken place, this can be determined by setting off the kernel compilation again, with

```
earth:/usr/src/linux # make bzImage
```

and watching out for corresponding error messages. But don't worry: errors do not occur very often when the kernel is being compiled!

If you are using the Bash shell you may enter:

```
earth:/usr/src/linux # make bzImage 2>&1 | tee kernel.out
```

for writing the compilation output to `kernel.out`. Using the Tcsh this command is invoked as:

```
earth:/usr/src/linux # make bzImage |& tee kernel.out
```

If you have configured parts of your kernel to be loaded as modules, you need to launch the compilation of modules after you have compiled the kernel. This may be achieved by typing:

```
earth:/usr/src/linux # make modules
```

10.6 Installing the Kernel

After having compiled a kernel, you have to make sure that it is available for use from now on.

If you use LILO, then this must also be reinstalled. In the simplest case, copy the the new kernel to `/boot/vmlinuz` (see Section 10.5 on the preceding page) and then run LILO to protect yourself from unpleasant surprises, however, it is recommended that you initially retain the old kernel (`/boot/vmlinuz.old`), so that you can still boot it if the new kernel does not function as expected:

```
earth:/usr/src/linux # cp /boot/vmlinuz /boot/vmlinuz.old
earth:/usr/src/linux # cp arch/i386/boot/bzImage /boot/vmlinuz
earth:/usr/src/linux # lilo
```

The Makefile target `make bzlilo` performs these three steps, by the way, in one go.

Now the compiled modules still need to be installed; by entering

```
earth:/usr/src/linux # make modules_install
```

you can copy these to the correct target directories in `/lib/modules/<Version>`. In doing so, the old modules will be overwritten with the same kernel version; you can, however, reinstall the original modules together with the kernel from the CDs.

Tip

You should make sure that modules whose functionality may now have been compiled directly into the kernel are removed from `/lib/modules/<version>`. Otherwise unexpected effects could occur! This is one reason why the inexperienced user is *strongly* advised not to compile the kernel himself.



In addition, enter a label **Linux.old** as boot image in your `/etc/lilo.conf` (also see Section 5.4.1 page 105) and rename the old kernel `/boot/vmlinuz.old`. Thus you can make sure that you can still boot using the previous kernel if it does not end up working with the new one. This option is described in detail in Chapter 5 page 99.

When you have adapted `/etc/lilo.conf` to your needs, you can enter:

```
earth:/usr/src/linux # lilo
```

If you boot Linux via DOS using `linux.bat` (`loadlin`), you have to copy the kernel to `/dos/loadlin/bzimage` (or to the directory where you have installed `loadlin`) for it to become active at the next boot.

If you start Linux via the Windows NT bootmanager you have to copy the LILO boot sector to the Windows NT start partition again (see Section 5.7.2 page 115 for details).

Furthermore, the file `/System.map` contains kernel symbols which are needed by the modules to launch kernel functions correctly. This file depends on the current kernel. Therefore, you should copy `/usr/src/linux/System.map` to the directory (`/boot`). If you create your kernel using `make bzlilo`, this is done for you automatically.

If you get an error message like "System.map does not match current kernel", then probably `System.map` has not been copied.

10.7 Creating a Boot Disk

If you want to create a boot disk with the new kernel, you can use the following command:

```
earth:/usr/src/linux # make bzdisk
```

10.8 Cleaning Your Hard Disk After Compilation

If you are low on hard disk space, you can regain some of the hard disk space used during compilation, with:

```
earth: # cd /usr/src/linux
earth:/usr/src/linux # make clean
```

If there is plenty of space and you plan to compile a new kernel from the same sources, you might skip the last step. A new compilation will then be faster, as only those parts of the system which must change based on your choices are recompiled.

11 Kernel Parameters

11.1 Drivers in the Kernel

There is a wide variety of PC hardware components. In order to be able to use this hardware properly, you need a “driver” with which the operating system (in Linux, the “kernel”), can access this hardware. In general there are two ways of integrating drivers into your system:

1. The drivers can be compiled directly into the kernel. Such a kernel (“in one piece”) is referred to as a *monolithic* kernel. Some drivers are only available in this form, justifying the need for monolithic kernels.
2. Drivers can be loaded on demand into the kernel, which is then referred to as a *modularized* kernel. This has the advantage that only those drivers are loaded which really are needed, and the kernel thus contains no unnecessary ballast. Our SuSE boot disk kernel makes use of modules to support most hardware configurations.

Regardless of whether the drivers are compiled directly into the kernel or are loaded as modules, it is still possible that a hardware component may not be recognized by the kernel. If this is the case, you have the option of specifying component attributes more exactly.

For monolithic kernels these parameters must be given at the boot prompt or by means of a boot loader.¹ Modular drivers get their parameters via `insmod` or `modprobe`, which load the module itself simultaneously.

You always have to give the parameters, when you boot your computer, as the kernel does not learn. Later, after the installation, kernel parameters can be set in the files `/etc/lilo.conf` or `/etc/modules.conf`, so that LILO or `modprobe` can interpret them automatically.

Unfortunately, the format for parameters that are compiled directly into the kernel is different from that for parameters that are loaded on demand. For this reason they are divided into two different groups in the following pages. A few modules now have unified parameters (e. g., CD-ROM drives), so that the same parameters can be given whether you are using modules or the boot prompt.

¹For this reason the parameters are also called LILO parameters, named after the time-honoured loader for X86 architectures

11.2 Some Tips

Before we present the list of parameters, here are some hints on recognizing hardware, adding parameters and booting with the SuSE boot disk:

- Most drivers can do an *autoprob*ing, e. g., test different standard addresses where this hardware can usually be found. While autoprob_{ing}, a driver might initialize an address that does not belong to it. This can cause the machine to hang.
- Some modules can be loaded successfully, even if their hardware is not installed. This applies mainly to 3Com network card drivers. You should try autoprob_{ing} first. Unused drivers can be easily removed. Hardware that is not automatically recognized must then be started using the appropriate parameters.
- Finally, there are some hardware components that have more than one driver (e. g., NCR 53C810, Ultrastor). As far as we know, there is no significant difference between the two Ultrastor drivers. The BSD driver for NCR53C810 also supports other NCR53C8xx products (e. g., 53C875), whereas the old NCR driver is the only one that recognizes CD writers. Use the one that suits your needs.

11.3 The Parameters

11.3.1 Notations and Meanings

Below you will see an alphabetical list of kernel parameters and the corresponding devices. The following parameters always have the same meaning:

<addr>	hexadecimal port address, (e. g., 0x300)
<irq>	interrupt the device will use (e. g., 7)
<dma>	DMA channel of the device (e. g., 1)
<Start address> ,	
<End address>	hexadecimal memory allocation for <i>shared memory</i>

Table 11.1: Frequently used variable names for kernel parameters

We are mainly concerned with the parameters that are required for a successful installation. But there are other parameters for special aims and purposes. Please keep in mind that upper case letters are significant.

A complete introduction to possible kernel parameters can be found after installation in the `BootPrompt-HOWTO`, in `/usr/share/doc/howto` as well as in the kernel sources in the file

`/usr/src/linux/Documentation/kernel-parameters.txt`.

11.3.2 Kernel Parameters at the Boot Prompt

The parameters listed in this section can only be entered to the kernel, for example, at the SYSLINUX (boot disk included), at the LILO prompt or via loadlin. If you plan to install one of these drivers as a module, please look at the parameters in Section 11.3.4 page 251.



Note

All parameters must be entered directly one after another, separated by commas. There should be no blanks (whitespace) between the parameters.

For example, to enter the parameters for the **aha1542** at the boot prompt you should enter (please keep in mind that you always have to precede the parameters with the boot configuration's name; on SuSE boot disks this is the configuration with the name `linux`):

```
linux aha1542=0x300
```

General boot parameters

By means of certain parameters the general behavior of the Linux kernel can be controlled.

- *Reboot mode* (shutting down Linux)

reboot=<mode>

Here, you have the following values for `<mode>`:

Variable	Values / Meaning
warm	warm reboot (no memory check)
cold	cold reboot (with memory check)
bios	BIOS reboot
hard	CPU crash reboot (triple fault)

Example: **reboot=cold**

This boots the machine after shutdown as if the reset button had been pressed.

- *Protecting memory segments* (reserve)

reserve=<start1>,<range1>,...,<startN>,<rangeN>

By means of this parameter, you can reserve memory space for hardware which does not like being autoprobed, and reacts with a system crash, for example.

Example: A “sensitive” network card can be protected from being autoprobed and initialized by entering:

reserve=0x330,32 ether=5,0x330,eth0

In this example, the network card has a 32 bit data bus which starts at address 0x330 and uses interrupt 5.

For the full description of network card parameters, see Section 11.3.2 page 247.

- *Passing root partitions*

root=<partition>

Variable	Values / Meaning
<partition>	e.g., /dev/hda1, /dev/sdb5

Example: **root=/dev/hda5**

This boots the kernel and tries to load the root partition from the first logical drive in the extended partition to the first (E)IDE hard disk.

- *Size of RAM (RAM)*

mem=<size>

You may enter the size of RAM in bytes, kilobytes or megabytes. Our example shows the different settings for 96 MB RAM.

Example:

mem=96M

mem=98304k

On some rare occasions it might happen that the mainboard cannot free the total amount of RAM. Please subtract the amount of memory the mainboard reserves (up to 512 K is considered normal). The exact amount can only be found by “trial and error”. Let’s assume that it is 512 KB for our example:

mem=5ff8000

When using a Pentium clone

mem=nopentium

this line might make it work.

- *Number of CPUs when using SMP*

maxcpus=<number>

Set the maximum number of CPUs to <number> when using an SMP kernel.

Example:

maxcpus=1

- *Serial console*

console=<device>,<options>

Set the device for the system console; this parameter can be used several times.

Examples:

console=ttyS1,9600

console=tty0

See /usr/src/linux/Documentation/serial-console.txt for more examples.

The SCSI controller and the SCSI subsystem

Most SCSI controllers can be configured via parameters.

- *SCSI Streamers*

st=<buffer>,<threshold>[,<max>]

Variable	Values / Meaning
<buffer>	Buffer size (Amount of 1 KB block)
<threshold>	Write threshold (Amount of 1 KB blocks)
<max>	Maximum buffer size (optional) (e. g. 2)

Example: **st**=1000,2000

- *Number of SCSI devices per ID*

max_scsi_luns=<number>

Variable	Values / Meaning
<number>	1..8

Example: If you want to use only the first LUN you need to set **max_scsi_luns**=1. This is typically used for imperfect CD changers where the amount of usable CDs equals the parameter **max_scsi_luns**=1.

- *Adaptec AHA-1520 / 1522 / 1510 / 1515 / 1505 SCSI host adapter*

aha152x=<addr>,<irq>,<id>[,<rec>[,<par>]]

Variable	Values / Meaning
<id> (SCSI ID of the host adapter)	0, 1
<rec> (reconnect)	0, 1
<par> (parity)	0, 1

This driver is used to run multiple low-cost SCSI controllers. All sound cards with a SCSI controller (except for Pro Audio Spectrum) can be accessed using the installed Adaptec chip.

For all non-original 152x, the fourth parameter (**RECONNECT**) seems to be necessary. For most types, it has to be set to '0'. Only the AHA2825 needs '1'.

Example: **aha152x**=0x300,10,7

- *Adaptec AHA-1540 / 1542 SCSI host adapter*

aha1542=<addr>[,<buson>,<busoff>[,<DMA speed>]]

Variable	Values / Meaning
<buson>	2..15
<busoff>	1..64
<DMA speed>	5, 6, 7, 8, 10

Example: **aha1542**=0x300

- *Adaptec AHA-274x / 284x / 294x Hostadapter*

aic7xxx=<modifier>[,<modifier>[, ...]]

Variable	Values / Meaning
<modifier>	extended activates translation of hard disk geometry no_reset disables reset of the SCSI bus at the host adapter initialization irq_trigger:<x> Only for Eisa systems 0 for flange triggered, 1 for lever driven verbose receive more messages reverse_scan If multiple cards are treated in the wrong order by the BIOS 7895_irq_hack:<x> -1 only for Tyan II Motherboards pci_parity:<x> if pci_parity is not used at all parity is even 0 no parity check 1 parity odd tag_info: " " " Queue management for performance enhancement, for experts, see kernel sources

Example: **aic7xxx=no_reset**,

if the machine hangs while resetting the SCSI bus.

parameters for aic7xxx based SCSI host adaptors are only necessary if they are faulty or malfunctioning.

From BIOS version 1.3 onwards the AHA-2940 AU causes no trouble. Updates may be obtained from Adaptec support.

The SCSI hostadapter Adaptec 2920 is driven by the Future Domain driver (see Section 11.3.2 on the facing page)

- *AdvanSys SCSI host adapter*

advansys=<addr1>,<addr2>,...,<addrN>

Example: **advansys=0x110,0x210**

This example tells the kernel to search for the AdvanSys host adapter at the given addresses.

- *AM53/79C974 SCSI host adapter*

AM53C974=<host-id>,<target-id>,<rate>,<offset>

Variable	Values / Meaning
<code><host-id></code>	the host adapter's SCSI ID (often 7)
<code><target-id></code>	the device's SCSI ID (often 0..7)
<code><rate></code>	3, 5, 10 MHz/s max. transfer rate
<code><offset></code>	transfer mode; 0 = asynchron

If the host adapter seems to “swallow” some packages, reduce the maximum transfer rate for this device (e. g., the first CD-ROM SCSI drive `/dev/scd0` with ID 5) on the SCSI bus with:

Example: `AM53C974=7,5,3,0`

- *BusLogic SCSI host adapter*

`BusLogic=<addr>`

`BusLogic=<probing>`

Example: `BusLogic=0x300`

Variable	Values / Meaning
<code><addr></code>	Address of the adapter, e. g. 0x300
<code><probing></code>	NoProbe No adaptor is probed NoProbeISA No ISA adapter is probed NoProbePCI No PCI adapter is probed NoSortPCI Order of multimaster adapter set by PCI BIOS MultiMasterFirst Multimaster before Flashpoint FlashPointFirst Flashpoint before Multimaster InhibitTargetInquiry For old devices that cause trouble with <code>scsi_luns > 0</code> TraceProbe outputs additional messages at initialization of the adapter TraceHardwareReset outputs additional messages at the adaptor hardware reset TraceConfiguration outputs additional messages at the adaptor configuration TraceErrors outputs error messages of the attached devices Debug outputs all

This host adapter understands even more parameters. These are for fine tuning the adaptor and are described in `/usr/src/linux/drivers/scsi/README.BusLogic`

- *Future Domain TMC-16x0 SCSI host adapter*

`fdomain=<addr>,<irq>[,<id>]`

Variable	Values / Meaning
<code><id></code>	SCSI ID of the host adapter 0..7

This driver also controls the Adaptec 2920.

Example: `fdomain=0x140,11,7`

- *Future Domain TMC-885/950 host adapter*

`tmc8xx=<addr>,<irq>`

Example: `tmc8xx=0xca000,5`

- *NCR 5380 SCSI host adapter family*

`ncr5380=<addr>,<irq>,<dma>`

Example: `ncr5380=0x340,10,3`

- *NCR 53c400 SCSI host adapter family*

`ncr53c400=<addr>,<irq>`

Example: `ncr53c400=0x350,5`

- *NCR 53c406a SCSI host adapter family*

`ncr53c406a=<addr>[,<irq>[,<fastpio>]]`

Variable	Values / Meaning
<code><fastpio></code>	0, if no fast PIO mode is required

Example: `ncr53c406a=0x330,10,0`

- *Seagate ST01/02 SCSI host adapter*

`st0x=<addr>,<irq>`

Example: `st0x=0xc8000,5`

- *Trantor T128/128F/228 SCSI host adapter*

`t128=<addr>,<irq>`

Example: `t128=0x340,10`

(E)IDE controllers and ATAPI devices

Numerous parameters are available to configure the (E)IDE controllers and the devices connected to them.

- *ATAPI CD-ROM* on the (E)IDE controller

`hd<x>=cdrom`

`hd<x>=serialize`

Variable	Values / Meaning
<code><x></code>	a, b, c, d

where:

Variable	Values / Meaning
a	master on 1st IDE controller
b	slave on 1st IDE controller
c	master on 2nd IDE controller
d	slave on 2nd IDE controller

Example: an ATAPI CD-ROM as master on the secondary IDE controller is set by `hdc=cdrom`.

- *Hard Disk*

`hd<x>=<cylinders>,<heads>,<sectors>[,<write>[,<irq>]]`

Variable	Values / Meaning
<x>	a, b, ..., h 1 to 8. Hard Disk
<cylinders>	number of cylinders
<heads>	number of heads
<sectors>	number of sectors
<write>	cylinders after which write compensation is used
<irq>	interrupt

If the BIOS is an old one, it is possible that the geometry of the hard drive is not recognized correctly. The correct parameters will then be passed on so that the kernel can still access the complete hard disk.

Example: `hdc=1050,32,64`

`hd<x>=<trouble>`

Variable	Values / Meaning
<x>	a, b, ..., h 1st to 8th hard disk
<trouble>	noprobe, if testing an existing hard disk causes problems none ignore CMOS entry and don't test nowerr ignore WREE_STAT-bit cdrom wrongly recognized as hard disk or not recognized at all, or will not boot autotune the fastest PIO mode is used slow adds a long break after each access. This makes it really slow, but it sometimes helps if it is the only option left

If a CD-ROM drive is not recognized reliably, specifying `<cdrom>` can register the device safely.

Example: `hdd=cdrom`

- *EIDE controller chipsets*

Some EIDE controllers have faulty chipsets or cause problems if the secondary controller is used.

Many of these chipsets are now supported in the kernel; this support needs to be specially activated, however, using a kernel parameter.

The following chipsets can be configured:

CMD 640	This chipset is found on many motherboards. Since it contains many bugs, the kernel offers a special support which recognizes the chip and bypasses the problems. Moreover, in some cases, use of the secondary controller is only possible by using this special support. In PCI systems, this chip is automatically recognized. For VLB systems, the following parameter is needed: ide0=cmd640_v1b .
RZ 1000	This chip is used on many motherboards that use the Neptune chipset and it is buggy. If support for this chip is activated, the system works a little slower but reliably. An additional activation with a kernel parameter is not required.
DTC-2278	Only activating this driver via ide0=dtc2278 makes it possible to use the secondary controller.
Holtek HT6560B	The following parameter is needed to activate the secondary controller: ide0=ht6560b .
QDI QD6580	If this driver is activated, it enables a higher speed: ide0=qd6580 .
UMC 8672	For activating the secondary controller, the following parameter is needed: ide0=umc8672 .
ALI M1439/M1445	For activating the secondary controller, the following parameter is needed: ide0=ali14xx .
PROMISE DC4030	For activating the secondary controller, the following parameter is needed: ide0=dc4030 . CD-ROMs and tapes on the secondary controller are not supported yet.

Table 11.2: Special EIDE Chipsets

If the chipset does not belong to the list of known faulty chipsets, and still is not recognized, the following parameters can be added instead:

```
ide<number>=<basis>[,<control>[,<irq>]]
```

Variable	Values / Meaning
<number>	adapter number, usually 0 or 1, but also 3 or 4
<basis>	base address of the adapter, usually 0x1f0, 0x170, 0x1e8 or 0x168
<control>	control register of the adapters, usually 0x3f6, 0x376, 0x3ee or 0x36e
<irq>	interrupt of the adapters, usually 14, 15, 11 or 10

If the chipset does not belong to the list of known faulty chipsets, but still causes problems, the following parameters can be added instead:

ide<number>=<tune>

Variable	Values / Meaning
<number>	adapter number, usually 0 or 1, but also 3 or 4
<tune>	autotune the highest possible PIO value is tried, not supported by all chipsets noautotune no improvement in speed serialize no time overlapping of operations with the next adapter

If the chipset does not belong to the list of known faulty chipsets, but the highest possible speed should still be achieved, the bus speed can be passed on; please refer to your motherboard manual.

idebus=<speed>

Other devices

- *Ethernet network cards*

ether=<irq>,<addr>[,<par1>[,<par2>...<par8>]],<Name>

The various parameters for <par1> to <par8> have different meanings for several drivers. Usually, only two parameters are needed, where the first is the start address and the second is the end address of *shared memory*. The first non-numerical argument is treated as the name.

<irq> the interrupt used; 0 for autoprobing
 <addr> port address; 0 for autoprobing
 <start> start address for shared memory. Some drivers use the 4 lowest bits for the debug level. The Lance uses them for its DMA channel.
 <end> end address for shared memory. The 3COM 3c503 driver uses this parameter to distinguish between internal and external transceivers.

Table 11.3: continued overleaf...

The Cabletron E21XX card uses the lowest four bits for selecting the media.

<Name> The interface's name (normally **eth0**)

Table 11.3: Variable names for Ethernet network cards

This parameter lets the kernel recognize more than one network card, since only the first card is searched by default. This can easily be done with:

```
ether=0,0,eth1
```

By passing 0 for both IRQ and address, the driver is told to launch *autoprob-
ing*, which means probing several addresses independently.

- *Floppy disk drives*

floppy=<drive>,<type>,cmos

Variable	Values / Meaning
<drive>	0, 1, 2, 3
<type>	0 – unknown or not recognized
	1 – 5 1/4" DD, 360 KB
	2 – 5 1/4" HD, 1.2 MB
	3 – 3 1/2" DD, 720 KB
	4 – 3 1/2" HD, 1.44 MB
	5 – 3 1/2" ED, 2.88 MB
	6 – 3 1/2" ED, 2.88 MB

floppy=<value>

For **<value>**, the following values may be set (Table 11.4 on the next page):

all_drives	more than two floppy disk drives
asus_pci	denies access to 3rd and 4th floppy disk drives
daring	use only with reliable controllers — enhances performance
0,daring	opposite of daring
<addr>,two_fdc	if the value <addr> is omitted while using a secondary floppy controller, a port address of 0x370 is set.
thinkpad	IBM Thinkpad machines
0,thinkpad	not a Thinkpad machine
omnibook	Omnibook computers
nodma	for Omnibook computers
dma	Standard

Table 11.4: continued overleaf...

nofifo	if a "Bus master arbitration error" occurs
fifo	Standard
0xX,fifo_depth	FIFO threshold standard 0xA
unexpected_interrupts	show warnings if something unexpected happens.
no_unexpected_interrupts	and...
L40SX	this value is the opposite of unexpected_interrupts .

Table 11.4: Kernel parameter values for **floppy**

- *Logitech bus mouse*

bmouse=<irq>

11.3.3 CD-ROM Drives on Proprietary Controllers

- *Aztech CDA268-01 CD-ROM*

aztcd=<addr>[,0x79]

The value 0x79 has only to be set when the firmware version is unknown.

Example: **aztcd**=0x320

- *Goldstar R420 CD-ROM drive*

gscd=<addr>

- *Mitsumi CD-ROM drive*

mcd=<addr>,<irq>[,<wait>]

Variable	Values / Meaning
<wait>	value for waiting time while starting

You can vary the parameter <wait> between 0 and 10 if the CD-ROM drive does not react fast enough on system requests ("timeout") and, therefore, does not find the root image while installing.

Example: **mcd**=0x300,10,5

- *Mitsumi CD-ROM drive (multisession)*

mcdx=<addr>,<irq>

Example: **mcd**=0x300,10

- *Mozart interface*

isp16=<addr>,<irq>,<dma>,<type>

Variable	Values / Meaning
<type>	Sanyo, Panasonic, Sony, Mitsumi

This driver is responsible for CD-ROM drives connected to the types ISP16, MAD16 or Mozart. The value of <type> follows the interface plug where the CD-ROM is connected to the sound card.

Example: `isp16=0x340,10,3,Sony`

- *Optics Storage 8000 AT CD-ROM drive*

`optcd=<addr>`

Example: `optcd=0x340`

- *Philips CM206 CD-ROM drives*

`cm206=<addr>,<irq>`

Example: `cm206=0x340,10`

- *Pro Audio Spectrum 16 – SCSI host adapter*

`pas16=<addr>,<irq>`

On the Pro Audio Spectrum 16 sound card there is a SCSI host adapter, the settings of which can be passed to the kernel with this parameter.

Example: `pas16=0x340,10`

If the card is not recognized, it can be operated without an interrupt. In this case the interrupt is set to 255.

Example: `pas16=0x340,255`

- *Sanyo CD-ROM Drive*

`sjcd=<addr>`

Example: `sjcd=0x340`

- *Sony CDU 31/33 A*

`cdu31a=<addr>,<irq>`

Autoprobing of this driver has been removed from the kernel. Thus, you have to set any parameters explicitly.

Example: `cdu31a=0x340,5`

If there is no interrupt assigned to the drive, and if you can only access it via polling mode, you have to enter 0 for the IRQ.

Example: `cdu31a=0x340,0`

As an example, if this drive is connected to a Pro Audio Spectrum card, the parameter could be as follows:

Example: `cdu31a=0x1f88,0,PAS`

- *Sony CDU 535*

`sonycd535=<addr>,<irq>`

Example: `sonycd535=0x340,10`

- *SoundBlaster Pro 16 MultiCD*

`sbpcd=<addr>,<type>`

Variable	Values / Meaning
<code><type></code>	LaserMate, SPEA, SoundBlaster

Example: `sbpcd=0x340,10`

The parallel port

- *Parallelport*

`parport=<addr0>,[<irq0> [parport=<addr1>,[<irq1>
[parport=<addr2>,[<irq2>]]]]]`

Variable	Values / Meaning
<code><addrX></code>	address
<code><irqX></code>	interrupt

Example: `parport=0x3bc parport=0x378,7 parport=0x278,auto`

- *Printer on the parallel port*

`lp=<parport0> [lp=<parport1> [lp=<parport2>]]`

Variable	Values / Meaning
<code><parportX></code>	parallel port

Example: `lp=parport0 lp=parport2`

11.3.4 modprobe Parameters

This section describes those parameters which can be loaded as modules. If you encounter difficulties while loading a driver (although you have entered the parameters), or if there is no section describing the parameter, then you may have to integrate this driver into a monolithic kernel.

Some drivers cannot be built as modules yet and some recognize your hardware properly only if they are compiled into the kernel. Despite this, we recommend you try the “module variant” first.

If a driver is loaded as a module, each and every variable used can be overwritten on the command line. There is, e. g., the variable `io` in the NE2000 driver which

specifies the used I/O range. For this, the correct command for loading this module is (see Section 10.2 page 230):

```
earth:/ # insmod ne io=0x300 irq=10
```

or better with modprobe:

```
earth:/ # modprobe ne io=0x300 irq=10
```

Be aware that there should be no spaces before or after the “equal”. Moreover, hexadecimal values have to be set in the given form (with a leading ‘0x’).



Note

If you want to enter more than one parameter they must be separated by spaces. This is the main difference to entering parameters at the LILO prompt, where no blanks should be used within the parameters for one driver.

The parameters you enter here can also be integrated into `/etc/modules.conf`. Here many parameters can be assigned to a particular module. This is done one line per module. The line should resemble the following:

```
options <module name> <parm1>=<value1> ...
```

where:

Variable	Values / Meaning
<module name>	name of the module without the extension .o
<parm1>	parameter #1
<value1>	value that is assigned to parameter #1

An entry for the NE2000 card might look like:

```
options ne io=0x300 irq=10
```

SCSI controllers and the SCSI subsystem

- *Adaptec AHA-1520 / 1522 / 1510 / 1515 / 1505 SCSI host adapter*

Name of module: `aha152x.o`

Variable	Values / Meaning
<io>	<addr>
<irq>	<irq>
<id>	SCSI-ID of the host adapter; by default, 7
<rec>	reconnect; 0, 1
<par>	parity; 0, 1
<sync>	synchronous mode of operation; 0, 1
<delay>	bus delay; default 100
<translat>	"translation of the hard disk geometry; 0, 1

As at the LILO prompt:

Example: `modprobe aha152x aha152x=0x340,10,7,1,1,0,0,0`

For a second host adapter:

Example: `modprobe aha152x aha152x1=0x140,12,7,1,1,0,0,0`

- *Adaptec AHA-1540 / 1542 SCSI host adapter*

Name of module: `aha1542.o`

Variable	Values / Meaning
<code><io></code>	<code><addr></code> ; default values are <code>0x330</code> and <code>0x334</code>
<code><buson></code>	2..15
<code><busoff></code>	1..64
<code><DMA speed></code>	5,6,7,8,10; the default is 5

At the boot prompt:

Example: `modprobe aha1542 aha1542=0x130`

- *Future Domain TMC-16x0 host adapter*

Name of module: `fdomain.o`

This driver also serves the SCSI host adapter Adaptec 2920, amongst others.

`fdomain=<addr>,<irq>[,<id>]`

Variable	Values / Meaning
<code><addr></code>	<code><addr></code>
<code><irq></code>	<code><irq></code>
<code><id></code>	<code><id></code>

As at the boot prompt:

Example: `modprobe fdomain fdomain=0x140,11,7`

- *NCR 5380 bzw. NCR 53C400 SCSI host adapter family*

Name of module: `g_NCR5380.o`

Parameter	Value
<code>ncr_addr</code>	<code><addr></code>
<code>ncr_irq</code>	<code><irq></code> ; switch off with 255
<code>ncr_dma</code>	<code><dma></code>
<code>ncr_5380</code>	1 for an NCR5380 adapter
<code>ncr_53c400</code>	1 for an NCR53C400 adapter

For an NCR5380 adapter:

Example: `modprobe g_NCR5380 ncr_irq=5 ncr_addr=0x350
ncr_5380=1`

For an NCR53C400 adapter with interrupts turned off:

Example: `modprobe g_NCR5380 ncr_irq=255 ncr_addr=0xc8000
ncr_53c400=1`

This driver serves the widely-used Trantor T130B SCSI host adapter

Network cards

- *3Com 3c501 / 3c503 / 3c505 / 3c507 network cards*

Name of module: 3c501.o, 3c503.o, 3c505.o, 3c507.o

Parameter	Value
io	<addr>
irq	<irq>

Example: `modprobe 3c509 io=0x300 irq=10`

- *3Com 3c509 / 3c579 network cards*

Name of module: 3c509.o

Parameter	Value
irq	<irq>
xcvr	0: internal; 1: external

Example: `modprobe 3c509 irq=10 xcvr=0`

- *3Com 3c515 network card*

Name of module: 3c515.o

Example: `modprobe 3c515`

- *3Com 3c59x / 3c90x network cards (“Vortex”/“Boomerang”)*

Name of module: 3c59x.o

For Compaq computers(PCI), you can try the following parameters:

Parameter	Value
compaq_ioaddr	<addr>
compaq_irq	<irq>
compaq_prod_id	<id>

Example: `modprobe 3c59x compaq_irq=10`

- *Allied Telesis AT1700 network card*

Name of module: at1700.o

Parameter	Value
io	<addr>
irq	<irq>

Example: `modprobe at1700 io=0x300 irq=10`

- *Cabletron E21xx network card*

Name of module: `e2100.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>
<code>mem</code>	<code><addr></code>
<code>xcvr</code>	0: internal; 1: external

Example: `modprobe e2100 io=0x300 irq=10 mem=0xd000 xcvr=0`

- *Digital DE425 / 434 / 435 / 450 / 500 network cards*

Name of module: `de4x5.o`

`io=0x<bus><device ID>`

Parameter	Value
<code>bus</code>	number of the PCI bus, usually 0
<code>device ID</code>	number des PCI device

With newer PCI BIOSes, these files are displayed when the machine boots, and in Linux they can be seen with the command

```
earth: # cat /proc/pci
```

Example: `modprobe de4x5 io=0x007`

- *Digital DEPCA / DE10x / DE20(012) / DE42, EtherWORKS network cards*

Name of module: `depca.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>
<code>mem</code>	<code><mem></code>
<code>adapter_name</code>	<code><name></code> e.g. DEPCA, de100, de101, de200, de201, de202, de210, de422

Example: `modprobe depca io=0x300 irq=10`

- *EtherWORKS 3 (DE203, DE204, DE205) network card*

Name of module: `ewrk3.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>

Example: `modprobe ewrk3 io=0x300 irq=10`

- *Intel EtherExpress 16 network card*

Name of module: `eexpress.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>

Example: `modprobe eexpress io=0x300 irq=10`

- *Intel EtherExpressPro network card*

Name of module: `eeepro.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>
<code>mem</code>	<code><addr></code>

Example: `modprobe eeepro io=0x300 irq=10 mem=0xd000`

- *Intel EtherExpressPro 100 network card*

Name of module: `eeepro100.o`

The Intel EtherExpressPro contains the chips i82557/i82558.

Parameter	Value
<code>options</code>	<code><duplex operation></code> <code><bus transceiver></code> 16 full--duplex 32 100 MBit operation only 64 10 MBit operation only

Example: `modprobe eeepro100 options=48`

This sets *full--duplex* and *100-MBit operation* simultaneously ($48 = 32 + 16$).

- *SMSC 83c170, "EPIC" family*

Name of module: `epic100.o`

Parameter	Value
<code>debug</code>	enable debugging 0..6
<code>options</code>	enable automatic detection: 0; for 10base2/BNC: 1
<code>full_duplex</code>	automatic detection: 0; full duplex: 1

Example: `modprobe epic100 debug=1 full_duplex=0,1`

This configures two cards, with full duplex being explicitly set on the second one.

You can read about additional configuration options at: <http://www.scyld.com/network/epic100.html>.

- *Fujitsu FMV-181/182/183/184 network cards*

Name of module: `fmv18x.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>

Example: `modprobe fmv18x io=0x300 irq=10`

- *HP PCLAN+ (27247B and 27252A) network card*

Name of module: `hp-plus.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>

Example: `modprobe hp-plus io=0x300 irq=10`

- *HP PCLAN (27245 / 27xxx)*

Name of module: `hp.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>

Example: `modprobe hp io=0x300 irq=10`

- *HP 10/100 VG-AnyLAN (ISA, EISA, PCI) network cards*

Name of module: `hp100.o`

Parameter	Value
<code>hp100_port</code>	<code><addr></code>

Example: `modprobe hp100 hp100_port=0x300`

- *ICL EtherTeam 16i / 32 network cards*

Name of module: `eth16i.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>

Example: `modprobe eth16i io=0x300 irq=10`

- *Novell NE2000 / NE1000 network cards*

Name of module: `ne.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>
<code>bad</code>	bad, only if the card is not recognized

Example: `modprobe ne io=0x300 irq=10`

- *NI6510 (AM7990 “lance” Chip) network card*

Name of module: `ni65.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>
<code>dma</code>	<code><dma></code>

Example: `modprobe ni65 io=0x300 irq=10`

- *SMC Ultra network card*

Name of module: `smc-ultra.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>

Example: `modprobe smc-ultra io=0x300 irq=10`

- *SMC 9194 network card*

Name of module: `smc9194.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>
<code>if_port</code>	<code><medium></code>

Variable	Values / Meaning
<code><medium></code>	0 auto 1 TP 2 AUI, 10base2

Example: `modprobe smc9194 io=0x300 irq=10 if_port=2`

- *Western Digital WD80x3 network card*

Name of module: `wd.o`

Parameter	Value
io	<addr>
irq	<irq>
mem	<mem>
mem_end	<mem_end>

Example: `modprobe wd io=0x300 irq=10`

- *IBM Tropic chipset Token Ring network card*

Name of module: `ibmtr.o`

Parameter	Value
io	<addr>
irq	<irq>
mem_start	<mem_start>

Example: `modprobe ibmtr io=0x300`

- *D-Link DE620 pocket adapter network card*

Name of module: `de620.o`

Parameter	Value
io	<addr>
irq	<irq>
bnc	1 if a BNC input/output
utp	1 if a UTP input/output
clone	1 if an equivalent device

Example: `modprobe de620 io=0x300 irq=10 bnc=1 utp=0`

Proprietary CD-ROM drives

The following parameters concern CD-ROM drives attached to special controllers. Whoever possesses one of these “oldies” will certainly know about it...

- *Aztech CDA268-01 CD-ROM drive*

Name of module: `aztcd.o`

Parameter	Value
aztcd	<addr>

Example: `modprobe aztcd aztcd=0x300`

- *Goldstar R420-CD-ROM drive*

Name of module: `gscd.o`

Parameter	Value
<code>gscd</code>	<code><addr></code>

Example: `modprobe gscd gscd=0x300`

- *Mitsumi CD-ROM drive*

Name of module: `mcd.o`

`mcd=<addr>,<irq>[,<wait>]`

Variable	Values / Meaning
<code><wait></code>	Delay at startup time

You can vary the `<wait>` from 0 to 10 if the CD-ROM drive does not react fast enough ("timeout"), thus probably not finding the root image during installation.

Example: `modprobe mcd mcd=0x300,10,5`

- *Mitsumi CD-ROM drive (Multisession)*

Name of module: `mcdx.o`

`mcdx=<addr>,<irq>[,<addr>,<irq>]`

For a number of drivers the settings `<addr>`, `<irq>` can be repeated up to 4 times.

- *Mozart sound card with port for CD-ROM drives*

Name of module: `isp16.o`

`isp16_cdrom_base=<addr>`

`isp16_cdrom_irq=<irq> isp16_cdrom_dma=<dma>`

`isp16_cdrom_type=<typ>`

Variable	Values / Meaning
<code><typ></code>	Sanyo, Panasonic, Sony, Mitsumi

This driver is not a real CDROM driver; it is only responsible for the interface configuration of the CD-ROM drive when connected to a ISP16, MAD16 or Mozart sound card. After this driver is loaded only the interface is loaded accordingly; the corresponding CDROM driver needs to be loaded as well. The value for the variable `<type>` is written on the plug which connects the CD-ROM cable to the sound card.

Example: `modprobe isp16 isp16_cdrom_base=0x300
isp16_cdrom_irq=10 isp16_cdrom_dma=1
isp16_cdrom_type=sony`

- *Optics Storage 8000 AT CD-ROM drive*

Name of module: `optcd.o`

Parameter	Value
<code>optcd</code>	<code><addr></code>

Example: `modprobe optcd optcd=0x300`

- *Philips CM206 CD-ROM drive*

Name of module: `cm206.o`

`cm206=<addr>,<irq>`

Example: `modprobe cm206 cm206=0x300,10`

- *Sanyo CD-ROM drive*

Name of module: `sjcd.o`

Parameter	Value
<code>sjcd</code>	<code><addr></code>

Example: `modprobe sjcd sjcd=0x300`

- *Sony CDU 31/33 A*

Name of module: `cdu31a.o`

Parameter	Value
<code>cdu31a_port</code>	<code><addr></code>
<code>cdu31a_irq</code>	<code><irq></code>

Example: `modprobe cdu31a cdu31a_port=0x300 cdu31a_irq=10`

- *Sony CDU 535*

Name of module: `sonycd535.o`

Parameter	Value
<code>sonycd535</code>	<code><addr></code>

Example: `modprobe sonycd535 sonycd535=0x300`

- *Soundblaster Pro 16 MultiCD*

Name of module: `sbpcd.o`

`sbpcd=<addr>,<typ>`

whereby `<type>` can have the following values:

Variable	Values / Meaning
0	LaserMate
1	SoundBlaster
2	SoundScape
3	Teac16bit

Example: `modprobe sbpcd sbpcd=0x300,0`

Sound cards and sound chipsets

Since the 2.2.xx kernel sound support has been resolved in modular form and can be controlled by numerous parameters.

- *AD1816 Chip*

Name of module: `ad1816.o`

The following cards, amongst others, are supported: Terratec Base 1, Terratec Base 64, HP Kayak, Acer FX-3D, SY-1816, Highscreen Sound-Boostar 32 Wave 3D

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>
<code>dma</code>	<code><dma></code>
<code>dma2</code>	<code><dma2></code>
<code>ad1816_clockfreq</code>	<code><clockfreq></code>

Example: `modprobe ad1816 io=0x530 irq=5 dma=1 dma2=3 ad1816_clockfreq=33000`

- *AD1848/CS4248 Chip (MSS)*

Name of module: `ad1848.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>
<code>dma</code>	<code><dma></code>
<code>dma2</code>	<code><dma2></code>
<code>type</code>	<code><cardtype></code>
<code>deskpro_x1</code>	<code><magic></code>

- *Generic OPLx driver*

Name of module: `adlib_card.o`

Parameter	Value
<code>io</code>	<code><addr></code>

Example: `modprobe adlib_card io=0x330`

- *Crystal 423x chipsets*

Name of module: `cs4232.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>
<code>dma</code>	<code><dma></code>
<code>dma2</code>	<code><dma2></code>

Example: `modprobe cs4232 io=0x530 irq=5 dma=1 dma2=3`

- *Ensoniq 1370 chipset*

Name of module: `es1370.o`

see also PCI64/128.

Parameter	Value
<code>joystick</code>	1: activate joystick
<code>lineout</code>	1: line-in in line-out switch over
<code>micz</code>	microphone impedancy

Example: `modprobe es1370 joystick=1 lineout=1`

- *Creative Ensoniq 1371 chipset*

Name of module: `es1371.o`

Refer also to the PCI64/128.

Parameter	Value
<code>joystick</code>	<code><addr></code>

Valid values for `<addr>` are `0x200`, `0x208`, `0x210` and `0x218`.

Example: `modprobe es1371 joystick=0x200`

- *Gravis Ultrasound*

Name of module: `gus.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>
<code>dma</code>	<code><dma></code>
<code>dma16</code>	<code><dma16></code>
<code>type</code>	
<code>gus16</code>	
<code>no_wave_dma</code>	
<code>dbl6</code>	

- *MAD16*

Name of module: `mad16.o`

Supports OPTi 82C928, OAK OTI-601D, OPTi 82C929, OPTi 82C930 and OPTi 82C924.

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>
<code>dma</code>	<code><dma></code>
<code>dma16</code>	<code><dma2></code>

Example: `modprobe mad16 io=0x530 irq=7 dma=0 dma16=1`

- *Turtle Beach Maui and Tropez*

Name of module: `maui.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>

Example: `modprobe maui io=0x530 irq=5`

- *MPU401*

Name of module: `mpu401.o`

Parameter	Value
<code>io</code>	<code><addr></code>

Example: `modprobe mpu401 io=0x330`

- *Turtle Beach MultiSound*

Name of module: `msnd.o`

Example: `modprobe msnd`

- *Turtle Beach Classic/Monterey/Tahiti*

Name of module: `msnd_classic.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>
<code>mem</code>	
<code>write_ndelay</code>	
<code>major</code>	
<code>fifosize</code>	
<code>calibrate_signal</code>	

Example: `modprobe io=0x290 irq=7 mem=0xd0000`

- *Turtle Beach Pinnacle/Fiji*

Name of module: `msnd_pinnacle.o`

As well as the parameters of the Turtle Beach Classic/Monterey/Tahiti, this card also understands:

Parameter	Value
<code>digital</code>	
<code>cfg</code>	
<code>reset</code>	
<code>mpu_io</code>	
<code>mpu_irq</code>	
<code>ide_io0</code>	
<code>ide_io1</code>	
<code>ide_irq</code>	
<code>joystick_io</code>	

Example: `modprobe msnd_pinnacle cfg=0x250 io=0x290 irq=5 mem=0xd0000`

- *OPL3*

Name of module: `opl3.o`

Parameter	Value
<code>io</code>	<code><addr></code>

Example: `modprobe io=0x388`

- *OPL3-SA1*

Name of module: `opl3sa.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>
<code>dma</code>	<code><dma></code>
<code>dma2</code>	<code><dma2></code>
<code>mpu_io</code>	<code><addr></code>
<code>mpu_irq</code>	<code><irq></code>

Example: `modprobe opl3sa io=0x530 irq=11 dma=0 dma2=1 mpu_io=0x330 mpu_irq=5`

- *YMF711, YMF715, YMF719, OPL3-SA2, OPL3-SA3, OPL3-SAx*

Name of module: `opl3sa2.o`

Parameter	Value
<code>io</code>	<addr>
<code>irq</code>	<irq>
<code>dma</code>	<dma>
<code>dma2</code>	<dma2>
<code>mss_io</code>	<addr>
<code>mpu_io</code>	<addr>

Example: `modprobe opl3sa2 io=0x370 irq=7 dma=0 dma2=3 mss_io=0x530 mpu_io=0x330`

- *Pro Audio Spectrum*

Name of module: `pas2.o`

Parameter	Value
<code>io</code>	<addr>
<code>irq</code>	<irq>
<code>dma</code>	<dma>
<code>dma16</code>	<dma16>
<code>sb_io</code>	<addr>
<code>sb_irq</code>	<irq>
<code>sb_dma</code>	<dma>
<code>sb_dma16</code>	<dma16>
<code>joystick</code>	
<code>symphony</code>	
<code>broken_bus_clock</code>	

- *Personal Sound System (ECHO ESC614)*

Name of module: `pss.o`

Parameter	Value
<code>pss_io</code>	<addr>
<code>mss_io</code>	<addr>
<code>mss_irq</code>	<irq>
<code>mss_dma</code>	<dma>
<code>mpu_io</code>	<addr>
<code>mpu_irq</code>	<irq>
<code>pss_mixer</code>	activate 1 or 0

- *Sound Blaster and Clones*

Name of module: `sb.o`

Parameter	Value
io	<addr>
irq	<irq>
dma	<dma>
dma16	<dma16>
mpu_io	<addr>
mad16	activate 1
trix	activate1
pas2	activate1
sm_games	activate1
acer	activate1, for Acer Notebooks
mwave_bug	activate1

Example: `modprobe sb io=0x220 irq=5 dma=1 dma16=5 mpu_io=0x330`

- *Aztech Sound Galaxy*

Name of module: `sgalaxy.o`

Parameter	Value
io	<addr>
irq	<irq>
dma	<dma>
dma2	<dma2>
sgbase	

- *S3 Sonic Vibes*

Name of module: `sonicvibes.o`

No parameters.

- *Ensoniq SoundScape*

Name of module: `sscape.o`

Parameter	Value
io	<addr>
irq	<irq>
dma	<dma>
mss	
mpu_io	<addr>
mpu_irq	<irq>
spea	1

- *MediaTrix AudioTrix Pro*

Name of module: `trix.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>
<code>dma</code>	<code><dma></code>
<code>dma2</code>	<code><dma2></code>
<code>sb_io</code>	<code><addr></code>
<code>sb_irq</code>	<code><irq></code>
<code>sb_dma</code>	<code><dma></code>
<code>mpu_io</code>	<code><addr></code>
<code>mpu_irq</code>	<code><irq></code>

- *UART401*

Name of module: `uart401.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>

Example: `modprobe io=0x330 irq=9`

- *UART6850*

Name of module: `uart6850.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>

- *Sound Blaster DSP chipsets*

Name of module: `v_midi.o`

No parameters.

- *Turtle Beach Maui, Tropez, Tropez Plus*

Name of module: `wavefront.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>

Example: `modprobe wavefront io=0x200 irq=9`

The parallel port

The operation of the parallel port is somewhat complex. After configuration, a complete subsystem is available; see Section 7.4.3 page 171 pp .

- *Parallel port*

Name of module: `parport.o`

Example: `modprobe parport`

- *Parallel port – architecture-specific*

Name of module: `parport_pc.o`

Parameter	Value
<code>io</code>	<code><addr></code>
<code>irq</code>	<code><irq></code>

In order to set up 3 Ports in a PC style, for example, whereby one has the address `0x3bc` without an IRQ, one has the address `0x378` with IRQ 7 and one has the address `0x278` with automatically detected IRQ, you should enter (all in one line!):

Example: `modprobe parport_pc io=0x3bc,0x378,0x278
irq=none,7,auto`

- *Printer driver*

Name of module: `lp.o`

Parameter	Value
<code>parport</code>	<code><port></code>

Example: `modprobe lp parport=0,2`

- *Parallel port IDE devices*

Name of module: `paride.o`

If you want to use chains of Parade devices on a *single* parallel port, then you must first load `parport.o` !

Example: `modprobe paride`

The parallel port should — if possible — be operated in “EPP mode”; please set this mode in the BIOS of your computer.

- *Parallel port IDE low-level protocol driver*

Name of module: `\variable{xxxx}.o`

protocol	device	
aten	ATEN EH-100	(HK)
bpck	Microsolutions backpack	(US)
comm	DataStor (old-type) "commuter" adapter	(TW)
dstr	DataStor EP-2000	(TW)
epat	Shuttle EPAT	(UK)
epia	Shuttle EPIA	(UK)
fit2	FIT TD-2000	(US)
fit3	FIT TD-3000	(US)
friq	Freecom IQ cable	(DE)
frpw	Freecom Power	(DE)
kbic	KingByte KBIC-951A and KBIC-971A	(TW)
ktti	KT Technology PHd adapter	(SG)
on20	OnSpec 90c20	(US)
on26	OnSpec 90c26	(US)

Example: `modprobe epat`

- *Parallel port IDE hard disk*

Name of module: `pd.o`

Parameter	Value
verbose	<value>

First load `parport` and the low-level driver

Example: `modprobe pd verbose=1`

- *Parallel port ATAPI CD-ROM*

Name of module: `pcd.o`

First load `parport` and the low-level driver

Example: `modprobe pcd`

- *Parallel port ATAPI floppy disk drive*

Name of module: `pf.o`

First load `parport` and the low-level driver

Example: `modprobe pf`

- *Parallel port ATAPI streamer drive*

Name of module: `pt.o`

First load `parport` and the low-level driver

Example: `modprobe pt`

- *Parallel port generic ATAPI device*

Name of module: `pg.o`

First load `parport` and the low-level driver

Example: `modprobe pg`

12 Updating the System and Package Management

12.1 Updating SuSE Linux

SuSE Linux offers the option of updating an existing system without reinstalling everything. It is important to distinguish between updating one or a few packages and updating the entire system to the latest distribution.

It is quite normal for software to “grow”. Thus we recommend that you have a look at how full the partitions are, with **df**, *before* updating! If you think there might be too little space available, consider making a backup and repartitioning your hard drives. There is no rule of thumb as to how much space you need in particular cases — disk space requirements depend on your existing partitioning scheme, the software installed, and on the version from which to update.



Note

It is helpful to read the file `README` or, in DOS/Windows `README.DOS`, on the CD. Additional changes are noted which may have been made *after* the manuals were printed.

12.1.1 Preparations

Before you begin an update, save the old configuration files to a separate medium, such as a streamer, removable drive, floppies, or ZIP drives, in case things go wrong. Normally these are the files in `/etc` and `/var/lib` (e.g., XDM). In addition, no harm would be done if you wrote the current user data in `/home` (the `HOME` directories) to a backup medium. A data backup should be carried out as `'root'`. Only `'root'` has the required permissions to read all local files.

Before you start the update process, write down the device name of your root partition, which can be determined by entering the command:

```
earth: # df /
```

Similarly, Output 12.1.1 on the following page shows that the root partition that to write down in our example is `/dev/hda7`.

As can be seen from the Output, `/dev/hda7` is mounted in the file system in `/`.

```
Filesystem      Size  Used Avail Use% Mounted on
/dev/hda7       3.0G  1.1G  1.7G  38% /
/dev/hda5       15M   2.4M   12M  17% /boot
shmfs          141M    0  141M   0% /dev/shm
/dev/hda1      1.9G  189M  1.7G  10% /dos
```

Output 12.1.1: File System Overview with `df -h`

Tip

Before updating *PostgreSQL* (package `postgres`), it is recommended to make a dump of the databases (see the man pages for `pg_dump`). This is only needed, however, if you used PostgreSQL *previously*.



Now you still need to decide if you want to update using the text-oriented YaST or with the graphical YaST2. The decision is made for you if you want to update from a SuSE Linux before version 6.0 — in this case you should use YaST and proceed as described in Section 12.1.3. >From version 6.0 onwards, you can update using YaST2.

12.1.2 Updating with YaST2

After the preparations mentioned in Section 12.1.1 on the preceding page, you should first boot.

Note

You can only carry out an update with YaST2 if you want to update a SuSE Linux 6.x system or newer. For older versions (e. g., SuSE Linux 5.3) YaST1 is the right method; see Section 12.1.3.



Start YaST2 and select 'Update existing system' from the menu. The module will go through several steps and YaST2 will determine which packages need to be updated. Decide this action for every single package, if desired. Using this module, you are not able to update the base system. To do this, boot from your installation disk, e. g., from CD.

12.1.3 Updating with YaST1

Updating the base system

Because central components of the system (such as libraries) must be replaced when updating the base system, this task cannot be performed during normal operations, i.e., from the currently running Linux system.

Start the update environment — normally using either the supplied “boot disk” or by entering **manually** at the kernel boot prompt, as described in detail in Section 2.1 page 7.

Essentially the following steps are necessary:

1. As soon as the kernel has booted from the boot disk or the update CD, linuxrc is started automatically.
2. In the main menu of linuxrc, in the menu ‘Settings’, choose the language, monitor, and keyboard. Press ‘OK’ when you are finished.
3. Now make sure that you load all the drivers needed by selecting ‘kernel modules’. The exact procedure is given in the linuxrc description in Section 16 page 294.
4. After this has been done, selecting ‘Installation / Startup system’ and ‘Startup installation’ leads to the selection of your source medium (see Section 16 page 296).
5. linuxrc then loads the installation environment and starts YaST automatically.

In YaST’s main menu, select ‘Update existing system’. YaST then tries to determine the root partition and displays the result. Here, select your root partition, as mentioned above (for example, /dev/sda3).

YaST then reads the existing “old” /etc/fstab and mounts the partitions it found in this file. Select ‘Continue’.

After returning to the main menu by pressing (**Esc**), choose ‘Updating your system’ (Section 3.2 page 43). Your old system is now analyzed by YaST and the results are displayed in a list.

After this, all the central configuration files of the system are updated. YaST creates backup files of those files that have changed since the last installation.

When the base system has been updated, you will find yourself in YaST’s special update mode, where you can decide which packages to update. In addition, old configuration files will be saved as .rpmorig or .rpmsave (see Section 12.3.2 page 279). This procedure is recorded in /var/adm/inst-log/installation-* and can be consulted later.

Updating the Rest of the System

When the base system has been updated, you will then be placed in YaST’s special update mode. Here, update the rest of your system as desired.

YaST then builds two lists. The first list shows those packages YaST has recognized and it thinks an update could be useful. The second list shows those packages where this is not so straightforward (perhaps the old package still works, but provides no information about its version).

Using these two lists, decide which packages should be updated. When you begin the update, all packages selected will be replaced by the new ones — all files will be saved that have been changed since the last installation.

After completing this task, proceed as with a first-time installation. One of the things you will need to do is select a new kernel. YaST will offer this option.



Tip

If you use `loadlin` for booting, copy the *new* kernel and possibly your `initrd` to the `loadlin` directory on your DOS partition.

If you do not want the system to be rebooted in its normal runlevel (see standard runlevel, Section 14.2 page 307), but would like YaST to complete the installation, enter the following at the LILO prompt:

```
NO_AUTO_SETUP=true
```

This parameter is especially useful if you encounter problems when booting again. These might occur if you access important parts of your system, for example, via PCMCIA-SCSI. To continue with the prepared configuration, the following steps could be taken:

1. Enter the following parameter at the LILO prompt:

```
NO_AUTO_SETUP=true
```
2. Change to user `'root'` and start YaST with `yast -nomenu` to complete its configuration tasks.
3. Start `/lib/YaST/bootsetup.conf` as `'root'`.

Experienced Linux users might not want to start in the default runlevel, but in the single user mode. This can be achieved by entering `single` at the LILO prompt.

12.1.4 Updating Single Packages

With SuSE Linux, update single packages whenever you want. When doing so, you must make sure that the system remains consistent: recommendations on updating can be found at <http://www.suse.de/en/support/download/updates/>.

In YaST's package list (see Section 3.5.3 page 63), move around as you please. If you select a package needed for the system to run, you will be warned by YaST. Such packages should be replaced only in update mode. For example, many packages contain *shared libraries* which could be in use when you want to run an update. Updating shared libraries on a running system will inevitably cause problems.

12.2 From Version to Version

In the following sections, we will list details of things which have been changed from one version to the next. This overview includes information on any basic settings that have been changed, configuration files that might have been moved,

or existing programs that have been visibly modified. Note, however, that the following list only includes those changes which directly affect users or administrators in performing their daily tasks. In other words, this list is by no means complete or exhaustive.

Also consult the Support Database — which can be found in the package `sdb_en`, series `doc`. Known problems and other abnormal features will be announced at: <http://sdb.suse.de/sdb/en/html>

Known problems and other abnormal features of each version will be announced on our web server when found. See the links given below. Important updates of individual packages can be found at <http://www.suse.de/en/support/download/updates/>.

12.2.1 From 6.4 to 7.0

Problems and Special Features:

<file:///usr/share/doc/sdb/de/html/bugs70.html>.

- All kernel modules (“drivers”) are included in the respective kernels installed (single kernel, multiprocessor kernel, etc.). This guarantees that the matching compiled modules are installed.
- The X Window System 4.0 no longer supports some very old graphics cards and a few very new ones. See Section 6.2 page 130. The setup program will be aware of this and switch to the preceding version 3.3.x, which is also included, accordingly.

The development environment of version 3.3.x should also be installed if you want to build packages in accordance with the “Linux Development Platform Specification”; see Section 13.3.1 page 286.

- `ypserv` from the package `ypserv` is no longer linked to the “tcp-wrapper” library, but instead uses `/var/yp/securenets`. After an update the settings from `/etc/hosts.allow` or from `/etc/hosts.deny` should be transferred to `/var/yp/securenets`.
- The portmapper is started via `/sbin/init.d/portmap` (since 7.1 via `/etc/init.d/portmap`) or with the command `rcportmap`; `/sbin/init.d/rpc` is now obsolete.
- In accordance with the FHS (*File system Hierarchy Standard*) (see Section 13.1.1 page 285), all architecture-dependent documentation can now be found in `/usr/share/doc` (previously `/usr/doc`).
- The package `bttv` is no longer included. The kernel modules are integrated into the respective kernel installed.
- The `gtk-XmHTML` library is available as a separate package `gxmhtml`; it has been moved from the package `gnlibs`.
- The package `gnadmin` has been discontinued; `logview` is now contained in the package `gnutils`.

- On the package `cron`: in accordance with the FHS, the cron tables are located in `/var/spool/cron/tabs`; see Section 13.3.2 page 286.
- On the package `postgres`: PostgreSQL and all of its components have been reorganized on the basis of the original packages. The package `pg_datab` with the initialization database is no longer needed; if needed, initialization is performed by the startup script.

12.2.2 From 7.0 to 7.1

Problems and Special Features:

<http://sdb.suse.de/sdb/de/html/bugs71.html>.

- *Kernel*: The “low-level” drivers for special EIDE chip sets are integrated into the standard kernel; `k_eide` or `eide` are no longer necessary as separate images. This means you no longer need to create your own boot disk if you require this driver.
- The meaning of *runlevel* has changed; see Table 14.1 page 308. The `init` scripts are now located in `/etc/init.d`; if you have already created your own scripts, you should make sure they are backed up before you update them.
- **DEFAULT_LANGUAGE**: New name for the former variable **LANGUAGE** in `/etc/rc.config`; see Section 13.9 page 303.
- The `/etc/resolv.conf` is now written directly from YaST as well as YaST2; no longer by SuSEconfig. Since version 7.2, this task is performed by the script `/sbin/modify_resolvconf`, included in package `aaa_base`.
- The limitation of the package names to 8 characters has been removed and many packages may now have more descriptive names. These new names will be taken into consideration in the update and the familiar programs will continue to be available to you automatically. If you are carrying out a reinstallation, you need to be careful that the packages are selected under their new names.
- Regarding the former package `ypclient`: This package is now divided into `ypbind` and `yp-tools` and the `init`-script is called `ypbind`.
- Regarding package `jade_dsl`: To avoid conflict with package `rzsz`, the command line tool `sx` is now called `s2x` and `sgml2xml`.

12.2.3 From 7.1 to 7.2

Problems and Special Features:

<http://sdb.suse.de/sdb/en/html/bugs72.html>.

- package `nkitb` no longer exists; the programs moved into several packages: `talk`, `rsh`, `finger`, `rwho`, `telnet`, etc., into the `-server` packages. For instance, the command `ping` can now be found in package `iputils`.

A number of programs are already “IPv6 ready”. For this reason you should ensure that DNS is correctly configured — otherwise it is possible you might have to wait for the DNS timeout for IPv6 queries. Also, in `/etc/nsswitch.conf`, replace the entry `dns` with `dns6`.

- Renaming of several packages (see on the preceding page); e. g.: series `sgm`: package `docbook-dsssl-stylesheets`, package `docbook_3` and package `docbook_4`.
- You have to install package `mod_php4` instead of package `mod_php` now.
- Emacs was split up into several packages:
 - Base package `emacs`.
 - In addition, you have to install package `emacs-x11`, which includes the program *with* X11 support.
 - package `emacs-nox` contains the program *without* X11 support.
 - package `emacs-info`: Online documentation in the info format.
 - package `emacs-el` contains uncompiled library files written in Emacs Lisp — not required for a runtime environment.
 - Numerous additional packages which can be installed, if required: package `emacs-auctex` (for \LaTeX); package `psgml` (for SGML/XML); package `gnuserv` (for a client-server environment), etc.
- Changes to follow the current FHS (see Section 13.1.1 page 285):
 - `/media` (was `/cdrom` and `/floppy`)
 Compatible links will be available for a transitional period.

12.2.4 From 7.2 to 7.3

Problems and Special Features:

<http://sdb.suse.de/sdb/en/html/bugs73.html>.

- Kernel 2.4 which is included in SuSE Linux has IDE-DMA turned on by default. The YaST2 module “System tuning” (package `yast2-config-tune`) as well as the back-end `idedma` have, therefore, become obsolete. If you have any problems related to DMA, you can always fall back on the ‘Installation -- Safe Settings’ boot option. Also see page 8.
- Apart from the `DISPLAYMANAGER` variable as defined in `rc.config`, the correct runlevel needs to be set in `/etc/inittab`. YaST2 is able to take care of this. Also see Section 4.6.2 page 87.
- Samba: The configuration files have been moved to the `/etc/samba` directory to for the sake of more clarity.

- MySQL now uses the TCP wrapper library for extended security. The new subpackage `mysql-max` offers all the new and enhanced features, which are automatically available after installation.
- To free up disk space (more than 30 MB), package `allman` is no longer offered with this distribution; however, the manual pages are still included in the respective packages.
- DocBook documents using certain “features” can be processed with `db2x.sh` (package `docbktls`). Note, however, that you should use the option `--strict` if you are required to produce compatible documents.

12.3 RPM — the Distribution Package Manager

In SuSE Linux, RPM (`rpm`), the “Red Hat Package Manager” serves as the package manager. It makes the powerful RPM database available to everyone: users, system administrators, and, of course, package builders. Detailed information on installed software can be queried from the database at any time.

`rpm` does the following:

- compiles software applications from “pristine” sources and packages them for installation
- installs, upgrades, and cleanly uninstalls software packaged in the RPM format
- supports queries, including dependencies, about packages and maintains the RPM database of installed packages

The reader is referred to the man page for `rpm` and the book *Maximum RPM*, Bailey, 1997, Red Hat, for more information on building RPM packages. The other capabilities of RPM are briefly described below.

Installable RPM archives are packed in a special binary format. These archives consist of the program files to be installed, as well as certain meta information which is used during the installation by `rpm` to configure the software package or stored in the RPM database for documentation purposes. RPM archives normally have the extension `.rpm`.

Tip



For a number of packages, the components needed for software development (libraries, headers, include files, etc.) have been put into separate packages. These development packages are only needed if you want to compile software *yourself*, for example, the most recent GNOME packages. These development packages can be identified by their name extension which is `-devel` (previously: `dev` or `d`): package `alsa-devel`, package `gimp-devel`, package `kdelibs-devel`, etc.

12.3.1 Verifying Package Authenticity

Since Version 7.1, the SuSE RPM packages have a GnuPG signature:

```
1024D/9C800ACA 2000-10-19 SuSE Package Signing Key <build@suse.de>
Key fingerprint = 79C1 79B2 E1C8 20C1 890F 9994 A84E DAE8 9C80 0ACA
```

With the command

```
earth:~ # rpm --verbose --checksig apache-1.3.12.rpm
```

verify the signature of an RPM package to determine whether it really originates from SuSE or from another trustworthy facility. This is especially recommended for update packages downloaded from the Internet. Our public package signature key normally resides in `/root/.gnupg/`.

12.3.2 Managing Packages: Install, Update, and Uninstall

Normally, RPM archive installation can be quickly taken care of:

```
earth: # rpm -i <package>.rpm
```

With this command, the package will be installed — but only if its dependency requirements are met and if it does not conflict with another package. With an error message, `rpm` requests those packages which need to be installed to fulfill dependencies. In the background, the RPM database ensures that no conflicts will arise — a specific file can only belong to one package. By choosing different options, force `rpm` to ignore these defaults, but be sure to know what you are doing, otherwise you risk compromising the integrity of the system and you might jeopardize the ability to update the system.

Use `-U` or `--upgrade` to update a package. This option will remove the files of the old version and immediately install the new files. `rpm` updates configuration files more cautiously:

- If a configuration file has *not* been changed by the system administrator, `rpm` will install the new version of the appropriate file. No action by the system administrator is required.
- If a configuration file was changed by the system administrator *before* the update, `rpm` will save the changed file with the extension `.rpmorig` or `.rpmsave` (backup file) and install the version from the new package, but only when the originally installed file and the newer version are different. If this is the case, compare the backup file (`.rpmorig` or `.rpmsave`) with the newly installed file and make your changes again in the new file. Afterwards, be sure to delete all `.rpmorig` and `.rpmsave` files to avoid problems with future updates.

The `-U` switch is *not* just an equivalent to uninstalling with the (`-e`) option and installing with the (`-i`) option. Use `-U` whenever possible.



Note

After every update, check all backup files created by `rpm`. These are your old configuration files. If necessary, incorporate your customizations from the backup files in the new configuration files. After this process, the files with the extensions `.rpmorig` and `.rpmsave` should be deleted.

To remove a package, enter the command:

```
earth: # rpm -e <package>
```

`rpm` will only delete the package if there are no unresolved dependencies. Theoretically, it is not possible to uninstall an old `libc` using `rpm` as long as another program still needs it to work properly — the RPM database guards against this.

12.3.3 RPM Queries

With the `-q` option, `rpm` initiates queries, making it possible to inspect an RPM archive (by adding the option `-p`) and also to query the RPM database of installed packages. Several switches are available to specify the type of information required (see Table 12.1).

<code>-i</code>	Package information
<code>-l</code>	File list
<code>-f <FILE></code>	Query a package owned by <code><FILE></code> (the full path must be specified with <code><FILE></code>)
<code>-s</code>	File list with status information (implies <code>-l</code>)
<code>-d</code>	list only documentation files (implies <code>-l</code>)
<code>-c</code>	list only configuration files (implies <code>-l</code>)
<code>--dump</code>	File list with complete details (to be used with <code>-l</code> , <code>-c</code> or <code>-d</code>)
<code>--provides</code>	List features of the package which another package can request with <code>--requires</code>
<code>--requires, -R</code>	Capabilities the package requires
<code>--scripts</code>	(Un)installation scripts (pre/post install/uninstall)

Table 12.1: The Most Important RPM Query Options (`-q [-p] ...<package>`)

For example, the command

```
earth: # rpm -q -i rpm
```

will display something like the following information:

```
Name       : rpm                               Relocations: (not relocateable)
Version    : 3.0.3                             Vendor: SuSE GmbH, Germany
Release    : 3                                 Build Date: Wed 21 Jul 1999 13:48:52
Install date: Wed 21 Jul 1999 18:44:33 MEST    Build Host: euler.suse.de
Group      : unsorted                          Source RPM: rpm-3.0.3-3.src.rpm
Size       : 5108780                           License: GPL
Packager   : feedback@suse.de
Summary    : RPM Package Manager
Description:
RPM Package Manager is the main tool for managing software packages
of the SuSE Linux distribution.
%rpm can be used to install and remove software packages; with rpm it's easy
%to update packages. rpm keep track of all these manipulations in a central
%database. This way it is possible to get an overview of all installed
%packages; rpm also supports database queries.
```

Option `-f` only works if you specify the complete file name with its full path. You can name as many file names as you want. For example,

```
rpm -q -f /bin/rpm /usr/bin/wget
```

will lead to the following result:

```
rpm-3.0.3-3
wget-1.5.3-55
% rpm-2.4.1-1
% wget-1.4.5-2
```

If you know only a part of the file name, you will have to use a shell script such as the following:

```
#!/bin/sh
for i in `rpm -q -a -l | grep $1 `; do
  echo you'll find "$i" in package":
  rpm -q -f $i
  echo ""
done
```

File 12.3.1: Script to Search for Packages

With the command

```
earth: # rpm -q --changelog rpm
```

have a precise list of information displayed (updates, configuration, modifications, etc.) on a specific package. This example is on the package *rpm*. With the help of the RPM database, verification checks can be made. These checks are initiated with the option `-V` (or `-y` or `--verify`). With this option, **rpm** will show all files in a package which have been changed since first being installed. **rpm** uses eight character symbols to give some hints about the kind of change (see Table 12.2 on the following page):

- 5 MD5 check sum
- S File size
- L Symbolic link

Table 12.2: continued overleaf...

T	Modification time
D	Major and minor device numbers
U	Owner
G	Group
M	Mode (permissions and file type)

Table 12.2: RPM Verify Options

In the case of configuration files, the letter `c` will be printed. For example, if you have changed `/etc/wgetrc` from the package `wget`, you may see:

```
earth: # rpm -V wget
S.5....T c /etc/wgetrc
```

The files of the RPM database are placed in `/var/lib/rpm`. If the partition `/usr` has a size of 1 GB, this database can occupy nearly 30 MB, especially after a complete update. If the database is much bigger than expected, it is useful to rebuild the database with the option `--rebuilddb`. Before doing this, make a backup of the old database. The `cron` script `cron.daily` makes daily copies of the database (packed with `gzip`) and stores them in `/var/adm/backup/rpmdb`. The number of copies is controlled by the variable `<MAX_RPMDB_BACKUPS>` (default: 5) in `/etc/rc.config`. The size of a single backup is approximately 3 MB. (This value is valid for a 1 GB `/usr` partition.) Take this space requirement into account when deciding how large you want the root partition to be. If `/var` has its own partition, you do not have to worry about this.

12.3.4 Installing and Compiling Source Packages

All source packages of SuSE Linux are located in the series `zq` (source packages) and carry an `.spm` extension (“Source RPMS”).

Tip

These packages can be handled in the same way as all other packages. The packages, however, will not be found in the RPM database (and are not marked with an `[i]` in YaST), as only “installed” software is listed.



The directories of `rpm` in `/usr/src/packages` must exist (if none of your own settings have been made, for example in `/etc/rpmdir`).

SOURCES this is for the original sources (`.tar.gz` files, etc.) and for distribution-specific adjustments (`.dif` files).

SPECS for the “spec” files, similar to a meta Makefile, which control the “build” process.

BUILD All the sources are unpacked, patched, and compiled in this directory.

RPMS This is where the completed “binary” packages are stored.

SRPMS here are the “source” RPMs.



Note

Do not experiment with essential system packages, such as package `libc`, package `rpm`, or package `nkit`. This could lead to a malfunctioning system!

When you install a source package from series `zq` with YaST, all the necessary components will be installed in `/usr/src/packages`: the sources and the adjustments in `SOURCES`, the relevant `.spec` file in `SPECS`¹. For our example, we will choose the `wget.spec` package. After you have installed the package with YaST, you should have the following files:

```
/usr/src/packages/SPECS/wget.spec
/usr/src/packages/SOURCES/wget-1.4.5.dif
/usr/src/packages/SOURCES/wget-1.4.5.tar.gz
```

`rpm -b <X> /usr/src/packages/SPECS/wget.spec` starts the compilation. Here `<X>` is a wildcard for different stages of the build process (see the output of `rpm --help` or the RPM documentation). Here is a short explanation:

- bp** Prepare sources in `/usr/src/packages/BUILD`: unpack and patch.
- bc** the same as **-bp**, but with additional compilation.
- bi** the same as **-bp**, but with additional installation of the built software. Caution: if the package does not support the BuildRoot feature, you might overwrite configuration files.
- bb** the same as **-bi**, but with the additional creation of the “binary” package. If the compile was successful, the binary should be in `/usr/src/packages/RPMS`.
- ba** the same as **-bb**, but with the additional creation of the “source RPM”. If the compilation was successful, the binary should be in `/usr/src/packages/SRPMS`.
- short-circuit** lets you skip specific steps.

This binary RPM may now be installed by invoking `rpm -i` or, even better, with `rpm -U` (to make it appear in the RPM database).

¹For “making packages” see [Bai97]. Further information can be obtained from the man page for `rpm`

12.3.5 Other Tools for Working with RPM Archives

Midnight Commander (**mc**) is able to “browse” RPM archives and to operate on parts of them. This tool works on an RPM package archive as if the archive were a regular file system. Using **mc**, view **HEADER** information with **(F3)** and copy parts of an archive with **(F5)**.

xrpm is a new graphical RPM manager, written in Python, which supports commands to FTP-accessed archives.

KDE can use the tool **krpm**, a graphical interface in the X Window System, for RPM management. **krpm** is currently in an early development stage.

Using the Alien (**alien**) Perl script, it is possible to convert or install an “alien” binary package. This tries to convert “old” TGZ archives to RPM before installing. This way, the RPM database can keep track of such a package after it has been installed. Beware: **alien** is still “alpha” software, according to its author.

Do not forget YaST.

13 Special Features of SuSE Linux

13.1 File System Hierarchy Standard (FHS) and the Linux Standard Base (LSB)

13.1.1 File System Hierarchy Standard (FHS)

SuSE Linux strives as far as possible to conform to the *File system Hierarchy Standard* (FHS, package `fhs`, series `doc`; see also <http://www.pathname.com/fhs/>). For this reason, it is sometimes necessary to move files or directories to their “correct” places in the file system.

13.1.2 Linux Standard Base (LSB)

SuSE supports the *Linux Standard Base* project. Current information on this can be found at: <http://www.linuxbase.org>. At this time, the LSB project exists in Version 0.3. Among other things, the system initialization is defined here. Refer to Chapter 14 page 307.

13.1.3 *teTeX* — \TeX in SuSE Linux

teTeX is put together according to the *TeX Directory Structure* (TDS) (see <ftp://ftp.dante.de/tex-archive/tds/>), without compromising the FHS.

13.2 Example Environments for FTP and HTTP

About FTP

To make it easier to set up an FTP server, the package `ftplib` package includes an example environment. This is installed in `/usr/local/ftp`.

About HTTP

Apache is the standard web server in SuSE Linux. Together with the installation of Apache, some example documents are made available in `/usr/local/httpd`. To set up your own web server, it is recommended that you enter your own `DocumentRoot` in `/etc/httpd/httpd.conf`.

13.3 Hints on Special Software Packages

13.3.1 The packages `xdevel` and `xdevel33`

The “Linux Development Platform Specification” specifies that XFree86 3.3.6 is to be used for building purposes, so that the software in question will run on as many platforms as possible. From SuSE Linux 7.0, however, XFree86 4.0 is the default. So that SuSE Linux can continue to meet the criteria of the specification, the package `xdevel33` is currently available.

To build a package with SuSE Linux in accordance with the “Linux Development Platform Specification”, uninstall the package `xdevel` and install the package `xdevel33` instead.

13.3.2 Package `cron`

The cron tables are now located in `/var/cron/tabs` (not, as previously, in `/var/lib/cron`). `/etc/crontab` serves as a system-wide cron table. Enter the name of the user who should run the command directly after the timetable (see File 13.3.1, here ‘root’ is entered). Package-specific tables, located in `/etc/cron.d`, have the same format — see the man page for `cron`.

```
1-59/5 * * * * root test -x /usr/sbin/atrun && /usr/sbin/atrun
```

File 13.3.1: Example of an Entry in `/etc/crontab`

`/etc/crontab` cannot be processed with `crontab -e`, but must be loaded directly into an editor, modified, then saved.

A number of packages install shell scripts to the directories `/etc/cron.hourly`, `/etc/cron.daily`, `/etc/cron.weekly`, and `/etc/cron.monthly`, whose instructions are controlled by `/usr/lib/cron/run-crons`. `/usr/lib/cron/run-crons` is run every 15 minutes from the main table (`/etc/crontab`). This guarantees that processes which may possibly have been neglected can be run at the proper time. Do not be surprised if, shortly after booting, the user ‘nobody’ turns up in the process tables and is highly active. This probably means that ‘nobody’ is just updating the locate database (see Section 14.6 page 322).

13.3.3 Package `curses`

On the CD, find the package `ncurses`. The corresponding libraries are named `libncurses.so.<xx>`. This means that some Makefiles have to be adapted to link `ncurses`. Packages of your own should be linked with the command `-lncurses` and never with `-lcurses`. For those who want to use `curses` anyway, try:

```
-I/usr/include/termcap -I/usr/include/curses
-L/usr/lib/termcap -L/usr/lib/curses
```

13.3.4 Sources for `uucp`

The sources of package `uucp` are included as a subpackage in the source RPM of Sendmail.

13.3.5 Man Pages

For some GNU applications (e. g., `tar`) the man pages are no longer maintained. They have been replaced by info files. Info (`info`) is GNU's hypertext system. Typing `info info` gives a starting help for using `info`. `info` can be launched via `emacs -f info` or on its own with `info`. The programs `tkinfo` and `xinfo` are easy to use or you can access the help system.

13.3.6 The Command `ulimit`

With the `ulimit` (*user limits*) command, it is possible to set limits for the use of system resources and to have these displayed. `ulimit` is especially useful for limiting the memory available for applications. Through this, an application can be prevented from using up too much memory on its own. This could bring the system to a standstill.

`ulimit` can be used with various options. To limit memory usage, use the options listed in Table 13.1.

- m maximum size of physical memory
- v maximum size of virtual memory (swap)
- s maximum size of the stack
- c maximum size of the core files
- a display of limits set

Table 13.1: `ulimit`: Setting Resources for the User

System-wide settings can be made in `/etc/profile`. There, creating core files must be enabled, needed by programmers for “debugging”. A normal user cannot increase the values specified in `/etc/profile` by the system administrator, but he can make special entries in his own `~/.bashrc`.

```
# Limits of physical memory:
ulimit -m 98304

# Limits of virtual memory:
ulimit -v 98304
```

File 13.3.2: `ulimit`: Settings in `~/.bashrc`

Details of memory must be specified in KB. For more detailed information take

a look at the man page for **bash**.

13.3.7 The File `/etc/resolv.conf`

The file `/etc/resolv.conf` will be updated by the script `/sbin/modify_resolvconf` exclusively, with no other program having permission to modify `/etc/resolv.conf` directly. Enforcing this rule is the only way to guarantee that the system's network configuration and the relevant files are kept in a consistent state.

13.4 Booting with the “initial ramdisk”

Describing the Problem

As soon as the Linux kernel has been booted and the root file system (`/`) mounted, programs can be run and further kernel modules can be integrated to provide additional functions.

To mount the root file system in the first place, certain conditions must be met. The kernel needs the corresponding drivers to access the device on which the root file system is located (especially SCSI drivers). The kernel must also contain the code needed to read the file system (`ext2`, `reiserfs`, `romfs`, etc.). It is also conceivable that the root file system is already encrypted. In this case, a password is needed to mount the file system.

If you just take a look at the problem of SCSI drivers, a number of different solutions are possible: the kernel could contain all imaginable drivers. This might be a problem because different drivers could conflict with each other. Also, the kernel will become very large because of this. Another possibility is to provide different kernels, each one containing just one or a few SCSI drivers. This method also has the problem that a large number of different kernels are required, a problem which is then increased by the differently optimized kernels (Pentium optimization, SMP).

The idea of loading the SCSI driver as a module leads to the general question which is answered by the concept of an *initial ramdisk*: creating a way of being able to run user space programs even before the root file system is mounted.

13.4.1 Concept of the Initial Ramdisk

The *initial ramdisk* (also called “initdisk” or “initrd”) solves precisely the problems described above. The Linux kernel provides an option of having a small file system loaded to a RAM disk and running programs there before the actual root file system is mounted. The loading of `initrd` is taken over by the boot loader (LILO, loadlin, etc.). All these boot loaders only need BIOS routines to load data from the boot medium. If the boot loader is able to load the kernel, it can also load the initial ramdisk. Special drivers are not required.

13.4.2 The Order of the Booting Process with `initrd`

The boot loader loads the kernel and the `initrd` to memory and starts the kernel. The boot loader informs the kernel that an `initrd` exists and where it is located in memory.

If the `initrd` was compressed (which is typically the case), the kernel decompresses the `initrd` and mounts it as a temporary root file system. A program called `linuxrc` is started on this in the `initrd`. This program can now do all the things necessary to mount the proper root file system. As soon as `linuxrc` finishes, the temporary `initrd` is unmounted and the boot process continues as normal with the mounting of the proper root file system. Mounting the `initrd` and running `linuxrc` can be seen as a short interlude during a normal boot process.

If the `initrd` cannot be unmounted, which is normally regarded as an error, the kernel will try and remount `initrd` to the directory `/initrd`. If the mount point `/initrd` does not exist, an error message will be displayed. In such a case, the system is fully functional, but the memory occupied by `initrd` cannot be released at all and is no longer available.

`linuxrc`

These are the only requirements for the program `linuxrc` in the `initrd`. It must have the special name `linuxrc` and it must be located in the root directory of the `initrd`. Apart from this, it only needs to be executable by the kernel. This means that `linuxrc` may be dynamically linked. In this case, the “shared libraries” in `/lib` must be completely available in `initrd`. `linuxrc` can also be a shell script. For this to work, a shell must exist in `/bin`. In short, `initrd` must contain a minimal Linux system that allows the program `linuxrc` to be carried out. When SuSE Linux is installed, a statically linked `linuxrc` is used, to keep `initrd` as small as possible (space on boot disks is very limited). `linuxrc` is carried out with `'root'` permissions.

The Real Root File System

As soon as `linuxrc` terminates, `initrd` is unmounted and discarded, the boot process carries on as normal, and the kernel mounts the real file system. What is mounted as the root file system can be influenced by `linuxrc`. It just needs to mount the `/proc` file system and write the value of the real root file system in numerical form to `/proc/sys/kernel/real-root-dev`.

13.4.3 Boot Loaders

Most boot loaders (above all, LILO, loadlin, and syslinux) can handle `initrd`. Individual boot loaders are given instructions on how to use `initrd` as follows:

1. LILO

Enter the following line in `/etc/lilo.conf`:

```
initrd=/boot/initrd
```

The file `/boot/initrd` is the *initial ramdisk*. It can be, but does not have to be, compressed.

2. loadlin.exe

run with:

```
C:> loadlin <kernelimage> \
    initrd=C:\loadlin\initrd <parameter>
```

3. syslinux

Enter the following line in `syslinux.cfg`:

```
append initrd=initrd <further parameters>
```

13.4.4 Using initrd in SuSE

Installing the System

The `initrd` has already been used for some time for the installation: here the user can load modules and make the entries necessary for an installation (above all, for the source medium). `linuxrc` then starts YaST, which carries out the installation. When YaST has finished, it tells `linuxrc` where the root file system of the freshly installed system is located. `linuxrc` writes this value to `/proc`, terminates, and informs the kernel to continue booting into the newly installed system.

For an installation of SuSE Linux, you are, from the very beginning, booting the system which is just being installed — somehow, quite clever. A real reboot after installation only takes place if the kernel does not match with the modules which were installed in the system. Since SuSE Linux only uses a kernel for uniprocessor systems, this will only happen if an SMP kernel was installed in the system with the corresponding modules. For this reason, an SMP kernel newly installed in the system must be rebooted for it to be able to use all modules.

Booting the Installed System

In the past, YaST has provided more than forty kernels for installing in the system, wherein the only basic difference in the kernels was that each of them contained a specific SCSI driver. This was necessary to be able to mount the root file system after booting. Further drivers could then be loaded afterwards as modules.

Because optimized kernels are now available, this concept is no longer feasible — by now, over one hundred kernel images would be needed.

This is why an `initrd` is used now, even to start the system normally. The way it is used is similar to that for an installation. The `linuxrc` used here, however, is

simply a shell script which just has the task of loading a given module. Typically, this is just one single module, namely the very SCSI driver which is needed to access the root file system.

Creating an `initrd`

An `initrd` is created by means of the script `mk_initrd`. In SuSE Linux, the modules to load are specified by the variable `INITRD_MODULES` in `/etc/rc.config`. After installation, this variable is automatically occupied by the correct values (the installation `linuxrc` knows which modules were loaded). Here it should be mentioned that the modules were loaded in exactly the same order in which they appear in `INITRD_MODULES`. This is especially important if a number of SCSI drivers are used, since otherwise the names of the hard disks would change. Strictly speaking, it would be sufficient just to load those drivers needed to access the root file system, because the automatic loading of additional SCSI drivers may cause problems (how should it be “triggered”, if hard disks hang on the second SCSI adapter), we load all SCSI drivers needed at the installation, by means of `initrd`.

The current `mk_initrd` checks if a SCSI driver is needed for the root file system. If you run `mk_initrd` on a system where `/` is found on EIDE hard disks, an `initrd` is not needed as the kernel used for SuSE Linux already contains the EIDE driver. As there are more and more special EIDE controllers coming onto the market, in the future it will probably be necessary, in such cases, to use an `initrd` to boot the installed system.



Note

Because the loading of the `initrd` with the boot loader runs in just the same way as loading the kernel itself (LILO notices in its `map` file the location of the files), LILO must be reinstalled after every change in `initrd`. After every `mk_initrd`, a `lilo` is also necessary.

13.4.5 Possible Difficulties — Self-Compiled Kernels

A self-compiled kernel can often lead to the following problems: out of habit, the SCSI driver is hard-linked to the kernel, but the existing `initrd` remains unchanged. When you boot, the following occurs: the kernel already contains the SCSI driver, the hardware is detected. `initrd`, however, now tries to load the driver again as a module. With some SCSI drivers, especially with the `aic7xxx`, this leads to the system blocking. Strictly speaking, this is a kernel error. An already existing driver should not be allowed to be loaded again as a module. The problem is already known from another context, however (serial drivers).

There are several solutions to the problem: either configure the driver as a module (then it will be correctly loaded in the `initrd`) or remove the entry for `initrd` from the file `/etc/lilo.conf`. An equivalent to the latter solution

is to remove the variable `INITRD_MODULES` then run `mk_initrd`, which then realizes that no `initrd` is needed.

13.4.6 Prospects

It is quite possible in the future that an `initrd` will be used for many more and much more sophisticated things than loading modules needed to access `/`.

- High-end EIDE drivers
- Root file system on RAID software (`linuxrc` sets up the `md` devices)
- Root file system on the LVM
- Root file system is encrypted (`linuxrc` asks for the password)
- Root file system on a SCSI hard disk on a PCMCIA adapter.

Further Information

```
/usr/src/linux/Documentation/ramdisk.txt
/usr/src/linux/Documentation/initrd.txt
manpage for initrd (man 4 initrd).
```

13.5 linuxrc

`linuxrc` is started during the boot of the kernel, usually as a prelude to a Linux system installation, before the “actual” booting commences (the kernel must first be properly configured, of course). This allows you to boot a small, modularized kernel and to load the few drivers you may need as modules from a floppy disk — in an emergency even from a second disk (`modules`).

`linuxrc` is your assistant for loading all relevant hardware drivers. You can also use `linuxrc` as a boot disk for an already installed system, for example, as a rescue disk. You can even start a totally independent RAM disk-based rescue system, for example, if something serious should happen to your hard disk or you have simply forgotten your `'root'` password. More about this below, in Section 13.6 page 297.

Main Menu

After you have selected the language, screen, and keyboard, find yourself in `linuxrc`'s main menu (see Figure 2.3 page 10). Start an installation with `'start installation / system'`. Whether you can go there directly depends mainly on your hardware.

If all components needed for installation have already been recognized by the kernel, you do not need to load additional drivers. This mainly applies to machines that only have (E)IDE adapters (and only (E)IDE hard disks and CD-ROM as well).

If there is a SCSI controller installed that is needed for installation,¹ load the corresponding SCSI module. The same applies if you want to install via an existing network. Here, the appropriate module has to be loaded first.

Furthermore, there are a lot of older CD-ROMs that are driven by proprietary controllers and which, therefore, need their own kernel modules. If PCMCIA devices are connected to a laptop, you need these modules as well.

System Information

If you are not sure about your hardware, the boot messages might help you.

In 'System information' (Figure 13.1), check a number of other things, besides kernel messages, such as the I/O addresses of PCI cards or the size of the main memory recognized by Linux.

You can check some system information in 'System information'. Here, check the used interrupts, I/O ports used, main memory, and recognized PCI devices as detected by Linux.

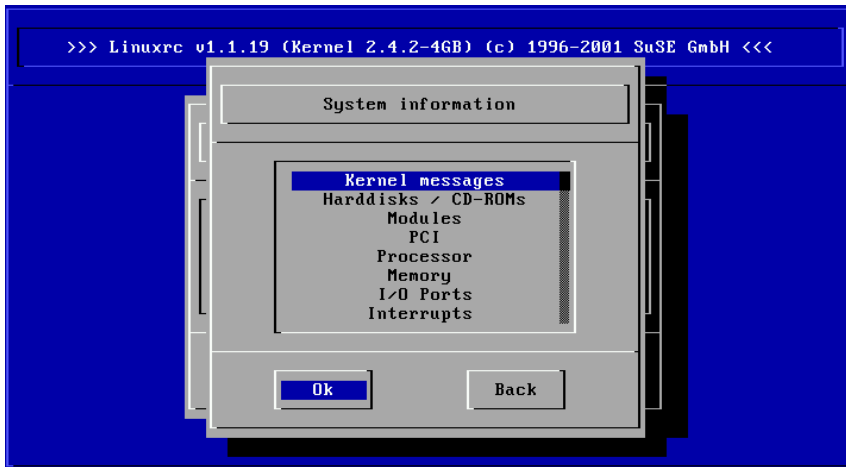


Figure 13.1: System Information

The next lines show how a hard disk and a CD-ROM connected to an (E)IDE controller announce their start. In this case, you do not need to load additional modules:

```
hda: ST32140A, 2015MB w/128kB Cache, LBA, CHS=1023/64/63
hdb: CD-ROM CDR-S1G, ATAPI CD-ROM drive
Partition check:
hda: hda1 hda2 hda3 < hda5 >
```

If you booted a kernel that already had a SCSI driver compiled, you do not need this SCSI driver as a module as well. Quite typical announcements when loading SCSI adapters and connected devices might look like this:

¹A controller with only a scanner connected to it is not required at start-up.

```

scsi : 1 host.
Started kswapd v 1.4.2.2
scsi0 : target 0 accepting period 100ns offset 8 10.00MHz FAST SCSI-II
scsi0 : setting target 0 to period 100ns offset 8 10.00MHz FAST SCSI-II
  Vendor: QUANTUM   Model: VP32210       Rev: 81H8
  Type:   Direct-Access          ANSI SCSI revision: 02
Detected scsi disk sda at scsi0, channel 0, id 0, lun 0
scsi0 : target 2 accepting period 236ns offset 8 4.23MHz synchronous SCSI
scsi0 : setting target 2 to period 248ns offset 8 4.03MHz synchronous SCSI
  Vendor: TOSHIBA   Model: CD-ROM XM-3401TA  Rev: 0283
  Type:   CD-ROM          ANSI SCSI revision: 02
scsi : detected 1 SCSI disk total.
SCSI device sda: hdwr sector= 512 bytes. Sectors= 4308352 [2103 MB] [2.1 GB]
Partition check:
sda: sda1 sda2 sda3 sda4 < sda5 sda6 sda7 sda8 >

```

Loading Modules

Select which kinds of modules you need. If you booted via disk, the corresponding data has to be read by `linuxrc` and displayed in a list.

If you have booted from CD or from DOS (via `loadlin`), these modules are already set in `linuxrc`. This saves tedious loading, but it needs additional memory. If your machine has less than 8 MB of RAM, boot from disk.

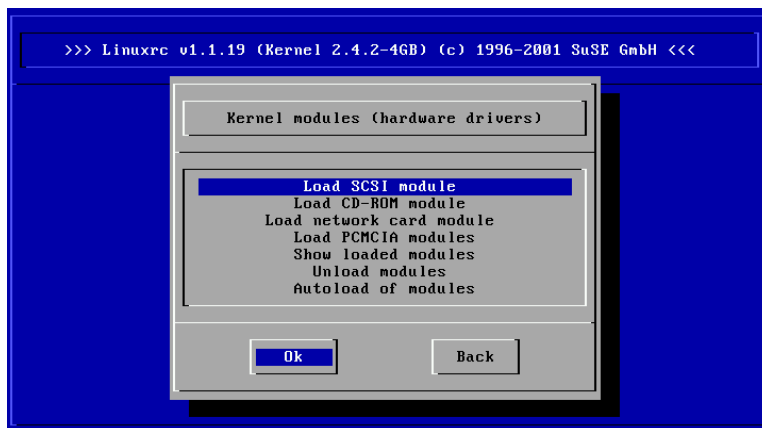


Figure 13.2: Load Modules

`linuxrc` offers a list of available drivers. On the left is the name of the module and, on the right, a short description of its usage.

For some components, there are a variety of drivers from which to choose (even newer alpha drivers).

Passing Parameters

When you have found a suitable driver, move to it with the cursor and press `(←)`. This opens a dialog box where you can add additional parameters for this module. More on module parameters can be found in Section 11.3.4 page 251.




Figure 13.3: Selection of SCSI Drivers




Figure 13.4: Entering Parameters for a Module to Load

We would like to point out that, in contrast to the LILO prompt, parameters for the same module have to be separated by blanks.

In most cases, it is not necessary to specify the hardware in detail. Most drivers find their components automatically. Most network cards and proprietary CD-ROM drives, however, need parameters. If in doubt, just try `(←)`.

Recognizing and initializing certain hardware can take some time. Switching to virtual console 4 (`(Alt) + (F4)`) lets you watch the kernel messages while loading. SCSI drivers need a while, as they have to wait for each device to load.

The messages are displayed by linuxrc so you can verify that everything ran smoothly. If it fails, the messages might give a hint as to why it failed.

Start Installation / System

Once you have set up hardware support via modules, switch to the ‘Start installation / system’ menu.




Figure 13.5: The linuxrc ‘Start’ Menu

From here (Figure 13.5), a number of procedures can be started: ‘Start installation’ (an update is also started from this item), ‘Boot installed system’ (the root partition must be known), ‘Start rescue system’ (refer to Section 13.6 on the facing page), and ‘Eject CD’.

‘Start Live-CD’ is only available if you booted a “LiveEval CD”. Download ISO images from the ftp server (`live-eval-<VERSION>`):

```
ftp://ftp.suse.com/pub/suse/i386/
```

Tip

The item ‘Start live-CD’ can be of great use if, for example, you want to test, without actually installing, if the computer in question or the notebook you might want to buy, is at all compatible with SuSE Linux — such a test ought to be possible in every modern PC shop, without any trouble.



For the installation (Figure 13.6 on the next page), choose various sources for the installation and similarly for the rescue system (see Figure 13.6 on the facing page).

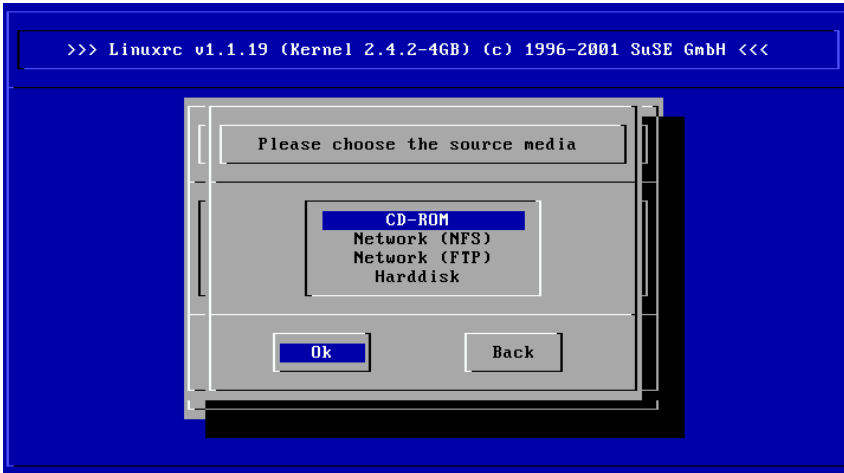


Figure 13.6: Selection of Source Media in linuxrc

13.6 The SuSE Rescue System

SuSE Linux contains several Linux rescue systems. With their help, reach your Linux partitions on the hard disks in emergencies “from the outside”: a boot disk or the “Rescue” System, loaded from disk, CD, the network, or the SuSE FTP server. There is also a bootable SuSE Linux CD (the “*LiveEval CD*”) which can be used as a rescue system.

Furthermore, the rescue system includes several help programs with which you can remedy large problems with inaccessible hard disks, misconfigured configuration files, or other similar problems.

Tip

Always invest some of your time in a boot and rescue disk, because the slight effort for creating and maintaining the disk is negligible in comparison to the time wasted and work involved when, in the case of emergency, you cannot access your system or your CD-ROM drive.



13.6.1 Preparations

For setting up your rescue system, you need two disks free of errors: one as future boot disk and the other for the compressed image of a small root file system. The image file `bootdisk` for booting the system and the file `rescue` for the root file system can be found on the first CD in the `/disks` directory.

There are three ways to set up the root file system on the disk:

- with YaST
- in a console using Linux commands


```
earth: # /sbin/badblocks -v /dev/fd0 1440
earth: # dd if=/media/cdrom/disks/rescue of=/dev/fd0 bs=18k
```
- at the DOS prompt (where the CD-ROM drive is Q:)


```
Q:\> cd \dosutils\rawrite
Q:\dosutils\rawrite> rawrite.exe
```

The rescue disk is currently based on libc5 (SuSE Linux 5.3), since it is possible in this to save some programs such as an editor, fdisk, or e2fsck to a disk.



Note

The rescue disk cannot be mounted, because it is not a file system. It only contains the compressed images of one. To read the file system, read the following paragraph.

If you want to read the image, you will need to decompress the image file then mount the decompressed image as user `'root'`. If your Linux kernel support the *loop device*, the procedure is as follows:

```
earth: # cp /media/cdrom/disks/rescue /root/rescue.gz
earth: # gunzip /root/rescue.gz
earth: # mount -t ext2 -o loop /root/rescue /mnt
```

13.6.2 Starting the Rescue System

The rescue system is launched using the SuSE boot disk or from your bootable SuSE Linux CD 1. It is required that the disk and CD-ROM drives are bootable. If necessary, you will need to change the boot series in the CMOS setup. Following are the steps for starting the rescue system:

1. Start your system with the SuSE boot disk or with the first SuSE Linux CD inserted in your CD-ROM drive.
2. Launch the entire system or, at the boot prompt, either enter `yast` or `manual`, where you can define which kernel modules should be loaded.
3. Make the respective settings for language, keyboard, and screen.
4. Select the item `'Installation/Start system'` in the main menu.
5. If you started with the boot disk, you should now insert the installation CD or the rescue disk with the compressed image of the rescue system.
6. In the menu `'Start installation/system'` select the item `'Start rescue system'` (see Figure 2.4 page 11) then specify the desired source medium (Figure 13.7 on the facing page). Subsequently, we will introduce a few tips on selection options:

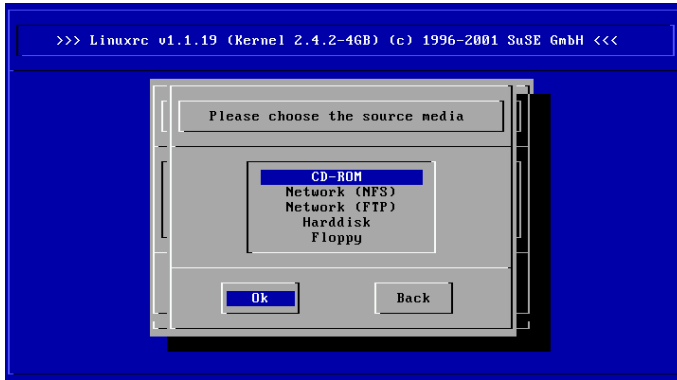


Figure 13.7: Source Medium for the Rescue System

'CD-ROM': When loading the rescue system, the path `/cdrom` is exported. This makes the installation from *this* CD possible.

Note

You now still need to enter the required values in `SuSEconfig` (Section 14.5 page 312).



'Network (NFS)': To start the `rescue` system via NFS from the network, you must have the driver for your network card already installed. Refer also to the general tips in Section 2.3.1 page 26).

'Network (FTP)': To start the `rescue` system via FTP from the network, you have to have your network card driver ready.

'hard disk': Load the `rescue` system from the hard disk.

'Floppy Disk': The `rescue` system can also be started from the floppy disk, especially if the computer only has a small amount of working memory.

Regardless of the medium chosen, the rescue system will be decompressed, loaded onto a RAM floppy disk as a new root file system, mounted, and started. Now it is ready for use.

13.6.3 Working with the Rescue System

The rescue system provides three virtual consoles on keys `(Alt) + (F1)` to `(Alt) + (F3)`. Here `'root'` may log in without a password. `(Alt) + (F4)` accesses the system console where you can view the kernel and syslog messages.

A shell and lots of other useful utilities (net tools), such as the `mount` program, can be found in the `/bin` directory. In `sbin`, find important file and network utilities for reviewing and repairing the file system (e.g., `e2fsck`).

Furthermore, this directory contains the most important binaries for system maintenance, such as `fdisk`, `mkfs`, `mkswap`, `mount`, `init`, and `shutdown`, as well as `ifconfig`, `route`, and `netstat` for maintaining the network.

An editor, `vi`, is located in `/usr/bin`. Also, tools like `grep`, `find`, and `less`, along with `telnet` are available.

Accessing Your Normal System

To mount your SuSE Linux system using the rescue system, use the mount point `/mnt`. You can also use or create another directory.

As an example, assume your normal system is put together according to the `/etc/fstab` shown in the example File 13.6.1.

<code>/dev/sdb5</code>	<code>swap</code>	<code>swap</code>	<code>defaults</code>	<code>0</code>	<code>0</code>
<code>/dev/sdb3</code>	<code>/</code>	<code>ext2</code>	<code>defaults</code>	<code>1</code>	<code>1</code>
<code>/dev/sdb6</code>	<code>/usr</code>	<code>ext2</code>	<code>defaults</code>	<code>1</code>	<code>2</code>

File 13.6.1: Example `/etc/fstab`

Caution

Pay attention to the order of steps outlined in the following section for mounting the various devices.



To access your entire system, mount it step-by-step in the `/mnt` directory using the following commands:

```
earth:/ # mount /dev/sdb3 /mnt
earth:/ # mount /dev/sdb6 /mnt/usr
```

Now you can access your entire system and, for example, correct mistakes in configuration files such as `/etc/fstab`, `/etc/passwd`, and `/etc/inittab`. The configuration files are now located in the `/mnt/etc` directory instead of in `/etc`.

To recover even completely lost partitions with the `fdisk` program by simply setting it up again, determine where on the hard disk the partitions were previously located and make a hardcopy printout of the `/etc/fstab` directory as well as the output of the command

```
earth: # fdisk -l /dev/<disk>
```

Instead of the `<disk>` variable, insert, in order, the device names of your hard disks, i.e., `hda`.

Repairing File Systems Damaged file systems are tricky problems for the rescue system. This could happen after an unscheduled shutdown caused by power failure or a system crash. Generally, file systems cannot be repaired on a running system. If you encounter really severe problems, you may not even be able to mount your root file system and have the system boot end in a "kernel

panic". Here, the only chance is to repair the system from the “outside” using a rescue system.

The SuSE Linux rescue system contains the utilities `e2fsck` and `dumpe2fs` (for diagnosis). These should remedy most problems.

In an emergency, man pages often are not available. That is why we have included them in this manual in Appendix A page 325.

Example:

If mounting a file system fails due to an *invalid* superblock, the `e2fsck` program would probably fail, too. If this were the case, your superblock may be corrupted, too. There are copies of the superblock located every 8192 blocks (8193, 16385, etc.). If your superblock is corrupted, try one of the copies instead. This is accomplished by entering the command:

```
earth: # e2fsck -f -b 8193 /dev/damaged_partition
```

The `-f` option forces the file system check and overrides `e2fsck`'s error so that — since the superblock copy is intact — everything is fine.

13.7 Virtual Consoles

Linux is a multiuser and multitasking system. You will learn to appreciate the advantages of these features, even on a stand-alone PC system.

In text mode, there are six virtual consoles available. Switch between them using `(Alt) + (F1)` to `(Alt) + (F6)`. The seventh console is reserved for X11.²

To switch to a console from X11 without leaving X11, use `(Ctrl) + (Alt) + (F1)` to `(Ctrl) + (Alt) + (F6)`. `(Alt) + (F7)` brings you back to X11.

13.8 Keyboard Layout

To standardize the keyboard mapping of programs, changes were made to the following files:

```
/usr/X11R6/lib/X11/Xmodmap
/etc/inputrc
/etc/skel/.exrc
/etc/skel/.less
/etc/skel/.lesskey
/etc/csh.cshrc
/etc/termcap
/usr/lib/terminfo/x/xterm
/usr/lib/X11/app-defaults/XTerm
/usr/share/emacs/<VERSION>/site-lisp/term/*.el
/usr/lib/joerc
```

These changes only affect applications that make use of `terminfo` entries or whose configuration files are changed directly (`vi`, `less`, etc.). Other non-SuSE applications should be adjusted to these defaults.

²You can assign more consoles via `/etc/inittab`.

Under X, the Compose key (“Multi_key”) can be accessed using the key combination **(Shift ↑) + (Ctrl (right))**; also see the corresponding entry in `/usr/X11R6/lib/X11/Xmodmap`.

13.9 Local Adjustments — I18N/L10N

SuSE Linux is, to a very large extent, internationalized and can be modified for local events in a flexible manner; in other words, internationalization (“I18N”) allows specific localizations (“L10N”). The abbreviations I18N and L10N are derived from the first and last letters of the words, and, in between, the number of letters omitted.

Settings are made via `LC_` variables defined in the file `/etc/rc.config.d/lang.rc.config`. This is less to do with “native language support” and more to do with the categories for News (Language), *Character Set*, *Sort Order*, *Time and Date*, *Numbers*, and *Money*. Each of these categories can be defined directly via its own variable or indirectly via an overlying variable in the file `rc.config` (see the man page for `locale`):

1. `RC_LC_MESSAGES`, `RC_LC_CTYPE`, `RC_LC_COLLATE`, `RC_LC_TIME`, `RC_LC_NUMERIC`, `RC_LC_MONETARY`: These variables are passed on to the shell without the `RC_` prefix and determine the above categories. The files concerned are listed below. `DEFAULT_LANGUAGE` is defined from `/etc/rc.config` directly and is *not* exported.

The current setting can be shown with the command `locale`.

2. `RC_LC_ALL`: This variable overwrites the values of the variables mentioned in item 1, if they have been set.
3. `RC_LANG`: If none of the above variables are set, this is the “fallback”. By default SuSE Linux only sets `RC_LANG`. This makes it easier for the user to enter his own values.
4. `ROOT_USES_LANG`: A *yes* or *no* variable. If it is set to *no*, ‘`root`’ always works in the POSIX environment.
5. `DEFAULT_LANGUAGE`: This variable has two functions. It defines the language in which YaST is started and, at the same time, is the fallback if `RC_LANG` is not set. It is set by YaST, depending on the language chosen in the language selection menu. To avoid any misunderstanding: `DEFAULT_LANGUAGE` is not passed on to the shell by `SuSEconfig` as `LANGUAGE`, but the value of this variable is only used to set `LANG`, if it is required.

`LANGUAGE` can also be used to define a “fallback chain” for language support. Read page 305 to see how this is done.

`DEFAULT_LANGUAGE` can be influenced by YaST via the item ‘Select language’. See Section 3.4.1 page 45. The other variables can be set via the `rc.config` editor. See Section 3.7.14 page 82.

The value of such a variable is put together from the “language code”, “country code”, “encoding,” and “modifier”. Individual details are connected by special characters:

```
LANG=<language>[_<COUNTRY>].Encoding[@Modifier]]
```

Some Examples

You should also use the language and country codes together. Language settings follow the standard ISO 639 (<http://www.indigo.ie/egt/standards/iso639/iso639-1-en.html> and <http://www.loc.gov/standards/iso639-2/>), country codes are listed in ISO 3166 (http://www.din.de/gremien/nas/nabd/iso3166ma/codlstpl/en_listpl.html). This only makes sense, however, if just those values are chosen for which usable description files can be found in `/usr/lib/locale`. Additional description files can be created using the command `localedef` from the files in `/usr/share/i18n`. A description file for `de_DE@euro.UTF-8` is thus created with:

```
earth:~ # localedef -i de_DE@euro -f UTF-8 de_DE@euro.UTF-8
```

LANG=de_DE.ISO-8859-1

This would set the variable to the German language, as used in Germany, and combine it with the character set `ISO-8859-1`. This character set does not support the Euro sign, but it will be useful sometimes for programs which have not been updated to support `ISO-8859-15`.

The string defining the charset (`ISO-8859-1` in our case) will then be evaluated by programs like Emacs.

LANG=de_DE@euro

This example shows how to use an option to request that the Euro symbol be used. Note that `de_DE@euro` is the default value for a standard installation in German.

LANG=de_DE.UTF-8

If you use a Unicode `xterm`, it is necessary to specify `UTF-8` as well. To achieve this, make a small shell script called `uxterm` to start `xterm` with `UTF-8` loaded each time; see File 13.9.1.

```
#!/bin/bash
export LANG=de_DE.UTF-8
xterm -fn \
'-Misc-Fixed-Medium-R-Normal--18-120-100-100-C-90-ISO10646-1' \
-T 'xterm UTF-8' $*
```

File 13.9.1: `uxterm` to Start a Unicode `xterm`

`SuSEconfig` reads the variables in `/etc/rc.config` and writes the necessary changes to `/etc/SuSEconfig/profile` and `/etc/SuSEconfig/csh.cshrc`. `/etc/SuSEconfig/profile` is read or “sourced” by `/etc/profile`. `/etc/SuSEconfig/csh.cshrc` is sourced by `/etc/csh.cshrc`. This makes the settings available system-wide.

Users can override the system settings by editing `~/.bashrc`. So, if the system settings are set to `en_US`, the user can, for example, switch to German system messages:

LC_MESSAGES=de_DE

13.9.1 Settings for Language Support

It should be pointed out that the files in the category *Messages* are, as a rule, only stored in the language directory (e. g., de) to have a fallback. If you set **LANG** to de_AT and the “message” file in /usr/share/locale/de_AT/LC_MESSAGES does not exist, it will fall back to /usr/share/locale/de/LC_MESSAGES.

A fallback chain can also be defined; e. g., for Breton → French or for Galician → Spanish → Portuguese:

```
LC_MESSAGES="br_FR:fr_FR"  
LC_MESSAGES="gl_ES:es_ES:pt_PT"
```

Further Information:

- *English-Howto* file: /usr/share/doc/howto/en/html/English-HOWTO.html
- Markus Kuhn, *UTF-8 and Unicode FAQ for Unix/Linux*, currently at: <http://www.cl.cam.ac.uk/~mgk25/unicode.html>.
- *Unicode-Howto*, by Bruno Haible
file: /usr/share/doc/howto/en/html/Unicode-HOWTO.html.

14 The SuSE Linux Boot Concept

Booting and initializing a UNIX system can challenge even an experienced system administrator. This chapter gives you a short overview of the SuSE Linux boot concept. The implementation at hand transcribes the *System Initialization* section of the LSB outline (Version 0.3); refer to Section 13.1.2 page 285 for more information on LSB.

The simple words "Uncompressing Linux..." signal that the kernel is taking control over your hardware. It checks and sets your console — or more precisely: the BIOS registers of graphics cards and output format — to read BIOS settings and to initialize basic hardware interfaces. Next, your drivers “probe” existing hardware and initialize it accordingly. After checking the partitions and mounting the root file system, the kernel starts `init` which “boots up” (Unix jargon) the main system with all its programs and configurations. The kernel will control the entire system, including hardware access and the CPU time programs may use.

14.1 The `init` Program

The program `init` is responsible for correctly initializing all system processes. Thus, it is the father of all processes in the entire system.

`init` takes a special role: `init` is started directly by the kernel and resists *signal 9*, which normally enables you to kill processes. All further programs are either started directly by `init` or by one of its “child” processes.

`init` is centrally configured via the `/etc/inittab` file. Here, the “runlevels” are defined (more about runlevels in the next section). It also specifies which services and daemons are available in each of the levels.

Depending on the entries in `/etc/inittab`, several scripts are invoked by `init` which, for reasons of clarity, all reside in the directory `/etc/init.d`.

The entire process of starting up the system (and shutting down, as well) is maintained by `init`. From this point of view, the kernel can be considered a “background process” whose task it is to maintain all other processes and to adjust CPU time and hardware access according to requests from other programs.

14.2 Runlevels

In Linux, *runlevels* define how the system is going to be started up. After booting, the system starts as defined in `/etc/inittab` in the line `initdefault`.

Usually this is 3 or 5 (see Table 14.1). An alternative to this is assigning a special runlevel at boot time (e. g., at the LILO prompt): the kernel passes any parameters which it doesn't need directly to `init`.

To change runlevels while the system is running, you can just enter `init` with the appropriate number. Only the super user is allowed to do this:

```
root@earth:/ > init 1
```

brings you to *single user mode* which is used for the maintenance and administration of your system. After finishing work in *S* mode, the system administrator can change the runlevel to 3 again by typing:

```
root@earth:/ > init 3
```

Now all essential programs are started and users can log in and work with the system. Table 14.1 below gives you an overview of available runlevels. Runlevel 2 should not be used on a system whose `/usr` partition is mounted via NFS:

Runlevel	Meaning
0	System halt
S	Single user mode; from boot prompt with US keyboard layout
1	Single user mode
2	Local multiuser without remote network (standard)
3	Full multiuser with network
4	Unused
5	Full multiuser mode with network and xdm
6	System reboot

Table 14.1: Valid Runlevels in Linux

You can halt the system using:

```
root@earth:/ > init 0
```

or cause it to reboot using:

```
root@earth:/ > init 6
```

If you have already installed and configured the X Window System properly (Section 6 page 129) and want users to log in via a graphical user interface, you can easily change the runlevel to 5. You should give it a try first by typing:

```
root@earth:/ > init 5
```

to see whether the system works as you expected.

Tip

With a damaged `/etc/inittab`, you can end up in a system which cannot be brought up properly. Therefore, be extremely careful while editing `/etc/inittab`! To repair the damage, you may try to enter `init=/bin/sh` at the LILO boot prompt for directly booting into a shell. This looks like: `boot: linux init=/bin/sh`



14.3 Changing Runlevels

Generally, there are a couple of things that happen if you change runlevels. First, *stop scripts* of the current runlevel are launched, closing down some programs which are essential for the current runlevel. Thereafter, *start scripts* of the new runlevel are started. Here, in most cases, a number of programs will be started.

To illustrate this, we will show you a change from runlevel 3 to 5:

- The administrator (`\root`) tells `init` to change runlevels:

```
root@earth:/ > init 5
```

- `init` now consults its configuration file (`/etc/inittab`) and realizes that it should start `/etc/init.d/rc` with the new runlevel as a parameter.
- Now `rc` calls all the stop scripts of the current runlevel, but only for those where there is no start script in the selected new runlevel. In our example, these are all the scripts which reside in `/sbin/init.d/rc3.d` (old runlevel was 3) and which start with a `\K`. The number following `\K` guarantees a certain order to start, as there are some dependencies which have to be taken into consideration.

Note

The names of the stop scripts always begin with `\K` for kill, and those for the start scripts with `\S` for start.



- The last thing to start are the start scripts of the new runlevel. These are (in our example) in `/sbin/init.d/rc5.d` and begin with an `\S`. The same procedure regarding the order in which they are started is applied here.

If you want to change to the same runlevel which you are already in, `init` only checks `/etc/inittab` for changes and starts the appropriate steps (e. g., for starting a `getty` on another interface).

14.4 Init Scripts

Scripts in `/etc/init.d` are divided into two sections:

- scripts which are executed directly by `init`. This only applies while booting and shutting down the system immediately (power failure or pressing `(Ctrl) + (Alt) + (Del)` by the user).
- scripts which are started indirectly by `init`. These are run when changing the runlevel and always call the master script `/etc/init.d/rc`, which guarantees the correct order of the relevant scripts.

All scripts are located in `/etc/init.d`. Scripts for changing the run level are also found there but are called via symbolic links from one of the subdirectories `/etc/init.d/rc0.d` to `/etc/init.d/rc6.d`. This is just for clarity reasons and avoids duplicate scripts (e. g., if they are used in several runlevels). Since every script can be executed as both a start and a stop script, these scripts have to “understand” the parameters “start” and “stop”. The scripts understand, in addition, the **restart**, **reload**, **force-reload** und **status** options; what these different options mean is elaborated in Table 14.2.

Option	Meaning
start	Starts service
stop	Stops service
restart	Stops service and restarts if service is already running; otherwise starts the service over again
reload	Read in configuration of service again without stopping and restarting it
force-reload	Read in configuration of the service again if the support supports this; in this way different than restart
status	show current status

Table 14.2: Summary of init Script Options

After leaving runlevel 3, `/etc/init.d/rc3.d/K40network` will be called up; `/etc/init.d/rc` calls up the `/etc/init.d/network` script by means of the **stop** parameter. After entering runlevel 5, the same script will be started, this time, however, with the **start** parameter.

Links in these runlevel–specific subdirectories simply serve to allocate the scripts to a certain runlevel.

Adding and removing the required links is done by the program `insserv` (or by the link `/usr/lib/lsh/install_initd`) when installing and uninstalling the package; refer to manpage for `insserv` (**man 8 insserv**)

Below, we give you a short introduction to the boot and stop scripts that are launched first (or last, respectively) as well as an explanation of the maintaining script.

- *boot*
Executed while starting the system directly using `init`. It is independent of the chosen runlevel and is only executed once. Here, file systems are checked, the kernel daemon is launched, some unnecessary files in `/var/lock` are deleted and the network is configured for the loopback device (if it has been selected in `/etc/rc.config`). Furthermore, the system time is set up with the `isapnp` tools and Plug and Play hardware configured (refer to Section 7.2 page 166).

If an error occurs while automatically checking and repairing the file system, the system administrator can intervene after first entering the root password.

The directory `/etc/init.d/boot.d` is assigned to this script. Every script, whose name starts with an `'S'` in this directory is executed while the system is booted. This is the right place for your own personal extensions, which should be executed only once.

Last to be executed is the script `boot.local`.

- *boot.local*
Here, you can enter additional commands to be executed at boot time before changing into a runlevel. It can be compared to `AUTOEXEC.BAT` on DOS systems.
- *boot.setup*
General settings that have to be performed while changing from *single user mode* to another runlevel.
Here, keyboard maps are loaded and the kernel daemon is started, which loads modules automatically.
- *halt*
This script is only executed while changing into runlevel 0 or 6. Here, it is executed either as `halt` or as `reboot`. Whether the system shuts down or reboots depends on how `halt` is called.
- *rc*
This script is of overriding importance whenever runlevels are changed. It calls the appropriate stop scripts of the current runlevel and the start scripts of the newly selected runlevel.


With this concept in mind, you can create your own scripts. A skeleton has been prepared in `/etc/init.d/skeleton`. The exact format is described in the LSB outline; this defines specifically the order of steps and in which levels the script has to be processed.

You may add your own scripts to this skeleton very easily. A template may be found in `/etc/init.d/skeleton`. To enable a script via `/etc/rc.config`, it is recommended that you create a `<START_>` variable in this file. Additional parameters should only be added if really needed (refer to the `/etc/init.d/gpm` script for reference).

Now you need to create the links in the corresponding `rc?.d` to your script to make sure it is launched when you change runlevels (see above Section 14.3 page 309 for script names, etc.). The manpage for `init.d` ([man 7 init.d](#)) gives you all the necessary technical background.

Caution

Please handle these scripts with the utmost care! A faulty script may hang your machine! See Section 14.2 page 307 if all else fails. . .



14.5 /etc/rc.config and SuSEconfig

Practically the entire configuration of SuSE Linux can be done via a central configuration file called `/etc/rc.config`. Here, a couple of environment variables are set which are (among others) checked by the init scripts. Each of the scripts in `/etc/init.d` executes `/etc/rc.config` as a first step, in order to read the values of those variables which apply to it.



Note

From SuSE Linux 6.0 packages with extensive adjustment options store their variables in individual files in the directory `/etc/rc.config.d`; this is the case, for example, with the package `sendmail` or the package `i4l` (ISDN).

To simplify things, however, we will usually just mention `/etc/rc.config`, even if the file has been moved somewhere else!

Moreover, a large number of configuration files can be generated from `/etc/rc.config`. This is the task of `/sbin/SuSEconfig`. If you change the network configuration, for example, the file `/etc/host.conf` will be regenerated, as it is dependent on the configuration you have made.

So, if you change `/etc/rc.config` manually, you should invoke `/sbin/SuSEconfig` afterwards to make sure all changes to the appropriate configuration files are made at the correct places. If you change the configuration with YaST, you don't have to bother. YaST automatically executes `/sbin/SuSEconfig` and updates your configuration files.

This concept enables you to make basic changes to your configuration without having to reboot the system. Since some changes are rather complex, some programs probably have to be restarted in order for the changes to take effect. This procedure is explained more fully in network configuration, where these programs are forced to be restarted using the commands:

```
earth:~ # rcnetwork stop
earth:~ # rcnetwork start
```

As you can see, you can easily start and stop init scripts by hand.

Generally, we recommend the following steps for configuring your system:

- Bring the system into *single user mode* (Runlevel 1):

```
earth:~ # init 1
```

- Change the configuration file `/etc/rc.config` as needed. This can be done using an editor of your choice or with YaST in 'Changing configuration file' (see Section 14.6 on the next page).
- Execute `/sbin/SuSEconfig` to make the changes take effect. If you have changed `/etc/rc.config` via YaST, this is done automatically.

- Bring your system back to the previous runlevel:

```
root@earth:~ > init 3
```

This procedure is mainly relevant if you have changed system-wide settings (such as network configuration). It is not necessary to go into *single user mode* for small changes, but it ensures that all relevant programs are correctly restarted.



Tip

To *categorically* disable the automatic configuration of SuSEconfig you need to set the variable `<ENABLE_SUSECONFIG>` in `/etc/rc.config` (please note Section 14.6). By using selected `rc.config` variables, you may partially disable the auto configuration.

14.6 The Variables in `/etc/rc.config`

In this section, we describe all the parameters of the system, including their default settings. If you don't use YaST to change `/etc/rc.config`, make sure you set "empty" parameters as two quotation marks (e. g., `KEYTABLE=""`) and that you surround parameters that contain a blank with quotation marks (parameters consisting of only one word do not have to be quoted). In our description, each parameter is given a value in order to make its settings as clear as possible:

Basic Settings

- **ENABLE_SUSECONFIG=yes**
With this entry you can disable SuSEconfig completely. Please don't contact our Installation Support if you have trouble configuring your system after disabling SuSEconfig ; -)
- **MAIL_REPORTS_TO=newbie**
SuSEconfig can mail reports (created by YaST or included in packages) to you. Here you can set the address. If you don't want this feature, simply set it to "".
- **GMT=-u**
If your hardware is set to GMT (*Greenwich Mean Time*), you should set this variable to `-u`.¹ Otherwise, leave it empty. This setting is relevant for automatic switching to summer or winter time.
- **TIMEZONE=CET**
The time zone where you live and where you operate your computer. This setting is also relevant for automatic switching to summer or winter time.

¹ `-u` is an abbreviation for *universal time*.

- **DEFAULT_LANGUAGE="english"**
Set by YaST when changes made using 'Define language' setting option (refer to Section 3.4.1 page 45). This value can also be classified as fallback if **RC_LANG** is not set; in this case, the value entered in `/usr/share/locale/locale.alias` will be resolved.
- **RC_LANG="en_US"**
Sets **LANG** for locale; this way, a default can be set for the local users. This value is comes in to play as long as no specific **RC_LANG_*** variables are used. The relevant `rc.config` variables are as follow: **RC_LANG_ALL** (This can overwrite **LANG_*** as well as **LANG!**), **RC_LANG_MESSAGES**, **RC_LANG_CTYPE**, **RC_LANG_MONETARY**, **RC_LANG_NUMERIC**, **RC_LANG_TIME** und **RC_LANG_COLLATE**. Refer to Section 13.9 page 303.
- **ROOT_USES_LANG="no"**
Should local settings also be used for 'root'?
- **ENABLE_SYSRQ="no"**
View inside of the kernel. Before implementing it, definitely read `/usr/src/linux/Documentation/sysrq.txt!`

Initializing Local Hardware (keyboard, modem, mouse, PCMCIA, etc.)

- **KEYTABLE=de-latin1-nodeadkeys**
Defines keymaps.
- **KBD_NUMLOCK=no**
(NumLock) on/off.
- **KBD_CAPSLOCK=no**
(CapsLock) on/off.
- **KBD_TTY="tty1 tty2 tty3 tty4 tty5 tty6"**
(NumLock) and (CapsLock) can be limited to certain TTYs; "" stands for all TTYs.
- **KBD_RATE=30**
Sets the automatic keyboard frequency. Possible values are from twice to 30 times per second. For this change to take effect, keyboard delay has to be set as well.
- **KBD_DELAY=250**
This is the delay whereafter automatic repetition of the pressed key commences. This value is in milliseconds but isn't very accurate. You have to assign **KBD_RATE** as well.
- **CONSOLE_FONT=mr.fnt**
This is the console font. Not all fonts support German umlauts. YaST provides a little window where you can test all fonts and choose the one you like best. Additional settings are: **CONSOLE_SCREENMAP**, **CONSOLE_UNICODEMAP** and **CONSOLE_MAGIC**.

- **MODEM=/dev/ttyS1**
Interface to which the modem is connected. YaST and SuSEconfig, in turn, create a link from `/dev/modem` to the device specified.
- **MOUSE=/dev/ttyS2**
Interface to which the mouse is connected. YaST and SuSEconfig, in turn, create a link from `/dev/mouse` to the device specified.
- **START_GPM=yes**
Set to `yes` to start mouse console support. This enables you to exchange text between consoles using the mouse. `gpm` can cause problems in connection with certain bus mice. If you encounter problems while starting X, you should disable `gpm`. The other alternative is to start `xdm`, since `gpm` is not started in runlevel 3.
- **GPM_PARAM="-t logi -m /dev/mouse"**
Initialization parameters for `gpm`. These are normally set via YaST.
- **PCMCIA=i82365**
This is for assigning the chipset: valid entries are: `i82365` and `tcic`. If the variable is set to `" "` the PCMCIA subsystem is not launched. Fine tuning is achieved via `PCMCIA_PCIC_OPTS` and `PCMCIA_CORE_OPTS`
- **START_ISAPNP=yes**
Initialize ISA P'n'P at boot time.
- **INITRD_MODULES="aic7xxx"**
The names of the modules which need to be loaded when the kernel is being booted, for example, to access the hard disk. Useful or necessary entries can be made during the installation or the update; see Section 13.4 page 288.
- **START_KERNELD=yes**
This variable determines whether the kernel daemon should be started automatically at boot time. This daemon is responsible for automatically loading kernel modules on demand. For the current standard kernel 2.2.xx, `kerneld` is *no longer* required. A short description of the module concept and functions of `kerneld` can be found in chapter Section 10.2 page 231.

Starting and Configuring Local Network and Other Services

- **START_LOOPBACK=yes**
Sets up a sort of “mini” network created by configuring the *Loopback* device. Since many programs rely on this functionality, it should be set.²
- **CHECK_ETC_HOSTS=yes**
SuSEconfig can do some checks and modifications to `/etc/hosts`.
- **BEAUTIFY_ETC_HOSTS=no**
If `/etc/hosts` is to be sorted.

²Your kernel needs to have been compiled with network support for this to work.

- **SETUPDUMMYDEV=no**
Determines whether the dummy network device should be set up: this is useful for non-permanent network connections (e. g., SLIP or PPP). If you have an ethernet card, it can be troublesome at times.
- **CREATE_HOSTCONF=yes**
SuSEconfig can create and check `/etc/host.conf`.
- **MODIFY_RESOLV_CONF_DYNAMICALLY=yes**
There are some services (“ppp”, “ipp”, “dhclient”, “pcmcia”), and “hotplug”) that have to change `/etc/resolv.conf` dynamically at certain times. For instance, if ppp/ipp establishes a connection, the peer might supply a list of name servers. Or PCMCIA needs to set the correct name server for the chosen configuration scheme.
If you don’t want these services to change `/etc/resolv.conf` at all, set this value to `no`. If unsure, leave it at the default (which is `yes`).
- **NETCONFIG=_0**
Specifies how many network cards (or other net devices) are installed. The text shows an example for one network card (they start with 0). For a system with two cards installed, it should resemble `NETCONFIG="_0 _1"`. For a system without networking, it should not be set.
- **IPADDR_0=193.141.17.202**
IP address of the first networking card.
- **NETDEV_0=eth0**
Name of the first network device (normally an Ethernet card, therefore, the example, `eth0`). Other possible settings are `str1` or `plip1`. If there is more than one card installed, additional cards are supplied with the variables `NETDEV_1` to `NETDEV_3`.
- **IFCONFIG_0="193.141.17.205 broadcast 193.141.17.255 netmask 255.255.255.192"**
Configuration command for the first networking device installed. These settings can easily be assigned using YaST. If you have more than one card installed, just enter the corresponding values in the appropriate variables.
- **CLOSE_CONNECTIONS=false**
If this variable is set to `true` *and* the system runs in “runlevel” 0 or 6, `/etc/init.d/route` sends a `SIGTERM` to all processes that own an open “remote tcp” or “udp” connection.
- **IP_DYNIP=no**
Switch on the “dynamic IP patch” when booting; if `yes`, the script `/etc/init.d/boot` enables this patch via an entry in the `/proc` file system.
- **IP_TCP_SYNCOOKIES=yes**
Activate “syn flood protection”.
- **IP_FORWARD=no**
If the machine should perform forwarding with two network cards,

`IP_FORWARD` should be set to `yes`; normally this is preferable, or required, for a “router” or for “masquerading”. The script `/etc/init.d/boot` enables “IP forwarding” by means of an entry in the `/proc` file system.

- **FQHOSTNAME=earth.cosmos.com**
Fully qualified hostname of your machine; fully qualified means the complete name, consisting of computer name and domain name.
- **ORGANIZATION="Gladstone Ganter Inc."**
This text appears in every news posting you send.
- **NNTPSERVER=sun**
Address of your news server. If you receive your news via UUCP and it is saved locally, you should enter `localhost`.
- **IRCSERVER=sun**
This is the place for your IRC server (*Internet Relay Chat*). Names of the servers should be separated by blanks.
- **START_INETD=yes**
Controls whether the `inetd` super daemon should be activated. This daemon reacts to calls from other hosts and starts (depending on the port) the appropriate service. You need it if you want to log in via `telnet` or `rlogin`. If you plan to use `xinetd` (see Section 19), set this to `no`.
- **START_XINETD=no**
Controls whether the `xinetd` super daemon should be activated (this is an enhanced `inetd`, see Section 19). If you plan to use this daemon, `START_INETD` should be set to `no`.
- **START_POSTFIX=no**
Activates the mail server postfix. The relevant variables are `POSTFIX_CREATECF`, `POSTFIX_RELAYHOST`, `POSTFIX_MASQUERADE_DOMAIN` and `POSTFIX_LOCALDOMAINS`.
- **SMTP=no**
Set to `yes` if a `sendmail` daemon should be activated. If you receive your e-mail exclusively via UUCP, you don’t need it, provided you call `sendmail -q` after each polling. (`rmail`, started by UUCP, just puts mail into a queue but doesn’t deliver it). If mail spool directories are mounted via NFS (on a network, for example) and the single host has got only outgoing mail, this could be set to `no` as well. The same applies with the use of *relay hosts*.
- **START_PORTMAP=no**
Determines whether to start the `portmapper` or not. You need the `portmapper` if you plan to use your host as an NFS server. Without this daemon, `rpc.mountd` and `rpc.nfsd` can’t run; this is why the `portmapper` is started when the variable is set to `no`, but `NFS_SERVER` is switched on! It is also necessary for NIS
- **NFS_SERVER=no**
If the host is going to be used as an NFS server, this variable has to be set to `yes`. This initializes the start of `rpc.nfsd` and `rpc.mountd`.

- **REEXPORT_NFS=no**
Set this variable to *yes* in order to re-export mounted NFS directories or NetWare volumes.
- **NFS_SERVER_UGID=yes**
If the daemon (`rcp.ugidd`) for the conversion of user and group ID's is to be started; this variable only takes effect if `<NFS_SERVER>` is set to *yes*.
- **USE_KERNEL_NFSD="no"**
If the package `knfsd` is installed, the kernel-based NFS daemon can be used. With this NFS daemon, locking is possible. See the corresponding variable, **USE_KERNEL_NFSD_NUMBER**.
- **START_AUTOFS=no**
This daemon enables you to automatically mount directories, via NFS as well as local directories, CD-ROM drives, floppy disks and more. The package `autofs`, series `n` must be installed and configured.
- **USE_NIS_FOR_AUTOFS=yes**
Whether to use NIS maps to fetch `autofs` configurations.
- **START_RWHOD=no**
Controls whether `rwhod` is started. Caution: the `rwhod` regularly sends "Broadcasts". If you use "on-demand" connection (ISDN or diald), this may cause connections to be made, resulting in extra traffic and costs!
- **START_ROUTED=no**
The route daemon is only necessary for dynamic routes (see the manpage for `routed` (`man routed`)). Caution, this service establishes a connection every 30 seconds. If the machine is connected to the Internet via dialup (for example, ISDN), then it makes *no* sense to set this variable to *yes*.
- **START_NSCD=yes**
NSCD Initializing "Name Service Caching Daemon" at boot time.
- **START_NAMED=no**
This controls whether the name daemon should be activated.
- **CREATE_YP_CONF=yes**
Set to *yes* if `SuSEconfig` should create the necessary files for YP. This depends on the next two entries as well. `SuSEconfig` also makes the appropriate changes to `/etc/passwd` and `/etc/group`.
- **YP_DOMAINNAME=cosmos.com**
YP domain name of hosts.
- **YP_SERVER=sun.cosmos.com**
Name of the NIS server.
- **USE_NIS_FOR_RESOLVING=no**
Use NIS for resolving host names.
- **START_CIPED=no**
Start CIPE daemon for an IPIP tunnel.

- **START_DHCPD=no**
Start server for DHCP *Dynamic Host Configuration Protocol*. For this the variables `DHCPD_INTERFACE`, `START_DHCRELAY` and `DHCRELAY_SERVERS` are used.
- **START_LDAP=no**
Start LDAP server.
- **START_RADIUSD=yes**
Start radius accounting and authentication service. This service is used by some ISPs for authentication of their users. See documentation in `/usr/share/doc/packages/radiusd`.
- **START_LPD=yes**
Start the lpd (“line printer”). Normally needed for printing.
- **START_NNTPD=yes**
Start nntpd if you want to access news via an NNTP site.
- **START_INN=no**
Start the INN news server.
- **START_ATD=yes**
Controls whether the AT daemon is activated. This daemon enables you to perform tasks periodically. In contrast to the cron daemon, this action is only performed once.
- **START_HTTPD=yes**
Controls whether the Apache http daemon should be activated.
- **START_HTTPSD=yes**
Specifies whether the Apache httpsd (“secure” webserver) with SSL and PHP3 should be started.
- **START_SQUID=no**
Controls whether the proxy server squid should be activated.
- **DOC_HOST=""**
If you want to use the central documentation server which contains the SuSE help system, enter the name of the host, e. g. `"sun.cosmos.com"`.
- **DOC_SERVER=no**
You should set this variable to `yes` on the documentation server. In `DOC_ALLOW` (see below) access to `http-rman` is set. Furthermore the index files for the http server have been rearranged: `http://`hostname-f`` instead of `http://localhost`.
- **DOC_ALLOW="LOCAL"**
List of machines (as patterns for `/etc/hosts.allow`) that are allowed to access the documentation server. This variable is read only if `DOC_SERVER` is set to `yes`. You could also enter a subdomain here (e. g. with `".cosmos.com"`).

- **HTTP_PROXY=""**
A couple of programs (e. g. lynx, arena, or wget) are capable of using proxy sites if this environment variable has been set.
SuSEconfig may set this in `/etc/SuSEconfig/*` (see SDB [file:/usr/share/doc/sdb/de/html/lynx_proxy.html](#)). Example:
`"http://proxy.provider.de:3128/"`.
- **FTP_PROXY=""**
FTP proxy. Example: `"http://proxy.provider.de:3128/"`.
- **GOPHER_PROXY=""**
Gopher proxy. Example: `"http://proxy.provider.de:3128/"`.
- **NO_PROXY=""**
This enables you to exclude (sub) domains from the proxy.
Example: `"www.me.de, do.main, localhost"`.
- **START_HYLAFAX=no**
Activates Hylafax. You will have to invoke `faxsetup` before setting this variable to `yes`.
- **START_SMB=no**
Start the samba server; a file and printer server for MS-Windows.
- **START_MARSNWE=no**
Activates the Novell server emulation.
- **START_SSHD=yes**
Start the “Secure Shell Daemon”; ensure before starting that a “host key” exists — see also the documentation in `/usr/share/doc/packages/ssh` as well as the man pages.
- **START_XNTPD=yes**
Controls whether the “Network Time Protocol (NTP) daemon” is activated (package `xntp`). It is configured via `/etc/ntp.conf`. Radio Controlled Clocks have addresses in the form `127.127.T.U`; in which `T` is the type of clock and `U` is the “unit number” which has a range from 0 to 3. Most clocks use a serial port or a special bus. The device file (Device) is normally only to the symbolic link `/dev/device-U` which refers to the actual hardware, where `U` has the value of the “unit number”. See also [/usr/share/doc/packages/xntp/html/refclock.htm](#). Example: You have a radio controlled clock attached to a serial port, you’ll also need the necessary Symlink. For this see `refclock.htm`. for the typical DCF77 receiver the “PARSE” driver is required:

```
## Type 8 Generic Reference Driver (PARSE)
## Address:      127.127.8.u
## Serial Port: /dev/refclock-u
```

So, if you choose, by an entry in `ntp.conf`, the value `server 127.127.8.0` you’ll also need the Symlink `/dev/refclock-0` to refer to `ttySx` — where `x` represents the serial port on which the Radio controlled clock is attached.

- **XNTPD_INITIAL_NTPDATE=""**
List of the NTP servers from which the time can be obtained before the local server is started; e. g. "sun.cosmos.com". If a number of servers are entered, they must be separated by a space. (package `xntp`). It is configured via `/etc/ntp.conf`.
- **DISPLAYMANAGER="kdm"**
Use this variable to define which display manager should be used for login. Possible entries are: `xdm` (the standard display manager of the X Window System), `kdm` (KDE's display manager), `gdm` (GNOME's display manager) or `wdm` (the "WINGs display manager").
- **KDM_SHUTDOWN=root**
Controls which user is allowed to shutdown the machine via `kdm` (reboot or shutdown). Possible values are: `root` ('root'), `all` (every user), `none`, and `local` (it may only be shutdown by users that logged in locally). If this is set to "", `root` is the default.
- **CONSOLE_SHUTDOWN=reboot**
Controls how `init` should react to `(Ctrl) + (Alt) + (Del)`. Possible values: `reboot` (the machine reboots), `halt` (the machine shuts down) and `ignore` (nothing happens). Default is `reboot`.
- **START_AXNET=no**
Applixware server for **Applixware** .
- **START_MYSQL=no**
Server for MySQL.
- **START_ADABAS=no**
Adabas server. The following variables belong to Adabas: `DBROOT`, `DBNAME`, `DBUSER` and `DBCNTROL` — see the respective comments in `rc.config`.
- **START_DB2=no**
Server for DB2.
- **START_ARKEIA=no**
Arkeia backup server.
- **START_ARGUS=no**
Argus server (network monitor).
- **ARGUS_INTERFACE=eth0** The interface Argus should listen to.
- **ARGUS_LOGFILE="/var/log/argus.log"**
The Argus logfile. This file might become quite large!
- **CRON=yes**
Sets the start and stop of the *cron daemon*. This daemon lets you start certain programs at a given time. It is only started in runlevels 2 and 3. It is highly recommended that you activate this daemon, especially if your computer is running all the time. An alternative or replacement is the AT daemon (see Section 19 page 319).



Note

There are a lot of options which require you to regularly run certain commands and programs. Therefore, the cron daemon should be activated on every system.

Local Maintenance

- **RUN_UPDATEDB=yes**
Set this to *yes* to have the locate database updated once per day via cron. locate is useful for quickly finding files. This tool may be fine tuned by a set of variables: **RUN_UPDATEDB_AS**, **UPDATEDB_NETPATHS**, **UPDATEDB_NETUSER**, and **UPDATEDB_PRUNEPATHS** (see comments in *rc.config*). — This updating may be carried out shortly after you have booted, if the computer is *not* switched on all the time; see Section 13.3.2 page 286.
- **REINIT_MANDB=yes**
If the man page data base should be renewed once a day by cron.daily.
- **CREATE_INFO_DIR=yes**
Set this to *yes* to have */usr/info/dir* created, which serves as a general index of all info pages. This is useful after installing a package which contains info pages. Keep in mind that perl needs to be installed for this to work.
- **CHECK_PERMISSIONS=set**
Controls check of file permissions according to */etc/permissions*. **set** corrects wrong entries, **warn** warns you, and **no** deactivates this feature.
- **PERMISSION_SECURITY="easy local"**
There are three levels of security: */etc/permissions.paranoid*, */etc/permissions.secure* and */etc/permissions.easy*. You may enter either **easy**, **secure** or **paranoid**. You can also create your own security levels; e.g. in */etc/permissions.local*. Then you may use **local** to activate your level.
- **RPMDDB_BACKUP_DIR=/var/adm/backup/rpmdb**
Controls where cron.daily should install its RPM database backup. If you want to disable this feature, set this variable to "".
- **MAX_RPMDDB_BACKUPS=5**
Maximum number of backup files for the RPM database.
- **DELETE_OLD_CORE=yes**
Core files are memory images of programs which have been killed due to a segmentation fault. These images are very useful for debugging purposes. If this is enabled, a regular search and deletion of old core files is launched.
- **MAX_DAYS_FOR_CORE=7**
Assigns how old core files should be before they are deleted.

- **MAX_DAYS_FOR_LOG_FILES=365**

If a log file (mainly in `/var/log`) reaches a certain size, it will be automatically compressed and archived. `'root'` is informed via e-mail. This parameter sets how long these files are to be kept on disk before being deleted automatically. If you set this value to 0, no compressing and no archiving will be done. Log files are written continuously and can grow to a considerable size! Compressed log files can be viewed anytime using `less`.

- **CLEAR_TMP_DIRS_AT_BOOTUP="no"**

Set this variable to `yes` to entirely remove (`rm -fr`) all files and subdirectories in the temporary directories, as defined through `TMP_DIRS_TO_CLEAR`.

Caution: If this variable is set to `yes`, the entries provided under `OWNER_TO_KEEP_IN_TMP` will *not* be taken into account, i. e. all files will be removed without exception!

- **MAX_DAYS_IN_TMP=30**

Selected directories (see `TMP_DIRS_TO_CLEAR` below) are checked daily to see whether they have been accessed during the selected time interval (set in days). Files which have not been modified will be deleted automatically.

- **TMP_DIRS_TO_CLEAR="/tmp /var/tmp"**

Enter here all directories which are to be searched for old files (see `MAX_DAYS_IN_TMP=30` above).

- **OWNER_TO_KEEP_IN_TMP="root bs"**

Files of system users given here should not be deleted, even if they have not been modified during the time specified.

- **CWD_IN_ROOT_PATH=no**

The current working current working directory in the path of user `'root'`.

- **ROOT_LOGIN_REMOTE=no**

If a login for `'root'` is to be allowed by `telnet` or `rlogin`, this variable should be set to `yes`. For security aspects concerning this variable, please refer to page [??](#).

- **SUSEWM_UPDATE=yes**

Controls whether `SuSEconfig` should adapt the system-wide configuration files according to the installed packages. You may fine tune this feature by means of the following variables: `SUSEWM_WM`, `SUSEWM_MWM`, `SUSEWM_XPM`, `SUSEWM_ADD` and `SUSEWM_COMPAT`.

A Manual Page of e2fsck

E2FSCK(8)

E2FSCK(8)

NAME

e2fsck - check a Linux second extended file system

SYNOPSIS

```
e2fsck [ -pacnyrdfvstFSV ] [ -b superblock ] [ -B block-size ] [ -l|-L bad_blocks_file ] [ -C fd ] device
```

DESCRIPTION

e2fsck is used to check a Linux second extended file system.

device is the special file corresponding to the device (e.g /dev/hdcl).

OPTIONS

-a This option does the same thing as the -p option. It is provided for backwards compatibility only; it is suggested that people use -p option whenever possible.

-b superblock

Instead of using the normal superblock, use an alternative superblock specified by superblock. This option is normally used when the primary superblock has been corrupted. The location of the backup superblock is dependent on the filesystem's blocksize. For filesystems with 1k blocksizes, a backup superblock can be found at block 8193; for filesystems with 2k blocksizes, at block 16384; and for 4k blocksizes, at block 32768. If an alternative superblock is specified and the filesystem is not opened read-only, e2fsck will make sure that the primary superblock is updated appropriately upon completion of the filesystem check.

-B blocksize

Normally, e2fsck will search for the superblock at various different block sizes in an attempt to find the appropriate block size. This search can be fooled in some cases. This option forces e2fsck to only try locating the superblock at a particular blocksize. If the superblock is not found, e2fsck will terminate with a fatal error.

-c

This option causes e2fsck to run the badblocks(8) program to find any blocks which are bad on the filesystem, and then marks them as bad by adding them to the bad block inode.

- C This option causes e2fsck to write completion information to the specified file descriptor so that the progress of the filesystem check can be monitored. This option is typically used by programs which are running e2fsck. If the file descriptor specified is 0, e2fsck will print a completion bar as it goes about its business. This requires that e2fsck is running on a video console or terminal.
- d Print debugging output (useless unless you are debugging e2fsck).
- f Force checking even if the file system seems clean.
- F Flush the filesystem device's buffer caches before beginning. Only really useful for doing e2fsck time trials.
- l filename
Add the blocks listed in the file specified by filename to the list of bad blocks. The format of this file is the same as the one generated by the badblocks(8) program.
- L filename
Set the bad blocks list to be the list of blocks specified by filename. (This option is the same as the -l option, except the bad blocks list is cleared before the blocks listed in the file are added to the bad blocks list.)
- n Open the filesystem read-only, and assume an answer of 'No' to all questions. Allows e2fsck to be used non-interactively. (Note: if the -c, -l, or -L options are specified in addition to the -n option, then the filesystem will be opened read-write, to permit the bad-blocks list to be updated. However, no other changes will be made to the filesystem.)
- p Automatically repair ("preen") the file system without any questions.
- r This option does nothing at all; it is provided only for backwards compatibility.
- s This option will byte-swap the filesystem so that it is using the normalized, standard byte-order (which is i386 or little endian). If the filesystem is already in the standard byte-order, e2fsck will take no action.
- S This option will byte-swap the filesystem, regardless of its current byte-order.
- t Print timing statistics for e2fsck. If this option is used twice, additional timing statistics are printed on a pass by pass basis.
- v Verbose mode.

- V Print version information and exit.
- y Assume an answer of 'Yes' to all questions; allows e2fsck to be used non-interactively.

EXIT CODE

The exit code returned by e2fsck is the sum of the following conditions:

- 0 - No errors
- 1 - File system errors corrected
- 2 - File system errors corrected, system should be rebooted if file system was mounted
- 4 - File system errors left uncorrected
- 8 - Operational error
- 16 - Usage or syntax error
- 128 - Shared library error

SIGNALS

The following signals have the following effect when sent to e2fsck.

SIGUSR1

This signal causes e2fsck to start displaying a completion bar. (See discussion of the -C option.)

SIGUSR2

This signal causes e2fsck to stop displaying a completion bar.

REPORTING BUGS

Almost any piece of software will have bugs. If you manage to find a filesystem which causes e2fsck to crash, or which e2fsck is unable to repair, please report it to the author.

Please include as much information as possible in your bug report. Ideally, include a complete transcript of the e2fsck run, so I can see exactly what error messages are displayed. If you have a writeable filesystem where the transcript can be stored, the script(1) program is a handy way to save the output of e2fsck to a file.

It is also useful to send the output of dumpe2fs(8). If a specific inode or inodes seems to be giving e2fsck trouble, try running the debugfs(8) command and send the output of the stat(1u) command run on the relevant inode(s). If the inode is a directory, the debugfs dump command will allow you to extract the contents of the directory inode, which can sent to me after being first run through uencode(1).

Always include the full version string which e2fsck displays when it is run, so I know which version you are running.

AUTHOR

This version of e2fsck was written by Theodore Ts'o <tytso@mit.edu>.

SEE ALSO

mke2fs(8), tune2fs(8), dumpe2fs(8), debugfs(8)

E2fsprogs version 1.19

July 2000

3

B The GNU General Public License

GNU General Public License

Copyright (C) 1989, 1991 Free Software Foundation, Inc.
59 Temple Place, Suite 330, Boston, MA 02111-1307, USA

Copyright (C) 1989, 1991 Free Software Foundation, Inc. 675 Mass Ave, Cambridge, MA 02139, USA Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Foreword

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the *GNU General Public License* is intended to guarantee your freedom to share and change free software — to make sure the software is free for all its users. This *General Public License* applies to most of the *Free Software Foundation's* software and to any other program whose authors commit to using it. (Some other *Free Software Foundation* software is covered by the *GNU Library General Public License* instead.) You can apply it to your programs, too.

When we speak of “*free*” software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU General, Public License

Terms and Conditions for Copying, Distribution and Modification

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this *General Public License*. The "Program", below, refers to any such program or work, and a *work based on the Program* means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

- b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
- c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

- a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, “complete source code” means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients’ exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The *Free Software Foundation* may publish revised and/or new versions of the *General Public License* from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and “any later version”, you have the option of following the terms and conditions either of that version or of any later version published by the *Free Software Foundation*. If the Program does not specify a version number of this License, you may choose any version ever published by the *Free Software Foundation*.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the *Free Software Foundation*, write to the *Free Software Foundation*; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

No Warranty

11. Because the program is licensed free of charge, there is no warranty for the program, to the extent permitted by applicable law. Except when otherwise stated in writing the copyright holders and/or other parties provide the program “as is” without warranty of any kind, either expressed or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. The entire risk as to the quality and performance of the program is with you. Should the program prove defective, you assume the cost of all necessary servicing, repair or correction.

12. In no event unless required by applicable law or agreed to in writing will any copyright holder, or any other party who may modify and/or redistribute the program as permitted above, be liable to you for damages, including any general, special, incidental or consequential damages arising out of the use or inability to use the program (including but not limited to loss of data or data being rendered inaccurate or losses sustained by you or third parties or a failure of the program to operate with any other programs), even if such holder or other party has been advised of the possibility of such damages.

End of Terms and Conditions

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the “copyright” line and a pointer to where the full notice is found.

one line to give the program’s name and a brief idea of what it does.
Copyright (C) 19yy *name of author*

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 675 Mass Ave, Cambridge, MA 02139, USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

```
Gnomovision version 69, Copyright (C) 19yy name of author
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details
type 'show w'. This is free software, and you are welcome to redistribute
it under certain conditions; type 'show c' for details.
```

The hypothetical commands `show w` and `show c` should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w` and `show c`; they could even be mouse-clicks or menu items — whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a “copyright disclaimer” for the program, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in the program
'Gnomovision' (which makes passes at compilers) written by James Hacker.
Unterschrift von Ty Coon, 1 April 1989
Ty Coon, President of Vice
```

This *General Public License* does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the *GNU Library General Public License* instead of this License.

C Important Keys

This chapter gives a short summary of some helpful and important key combinations.

`←`

Executes a command on the command line; otherwise, a linefeed.

`Alt` + `F1`–`F6`

Change to another console (in text mode).

`Ctrl` + `Alt` + `F1`–`F6`

In X11 and Dosemu: change to one of the text consoles.

`Alt` + `F7`

Switch from a text console back to X11.

`Ctrl` + `Alt` + `Backspace`

Exit X11.

`Shift` ↑ + `Page` ↑

Shows previous screens on console. Works until consoles are switched.

`Shift` ↑ + `Page` ↓

Opposite of `Shift` ↑ + `Page` ↑.

↑

Displays the previous commands on the command line.

↓

Just the opposite of ↑ (in the shell).

`Shift` ↑ + `Alt` + `Cursor`

X11: Change to the next virtual desktop (in the direction of the cursor). The number of virtual desktops is set in `.fvwm[2]rc` with the parameter **DeskTopSize**.

`Ctrl` + `Alt` + `±(gray)`

Changes the screen resolution in the order of the monitor settings made in `/etc/XF86Config`.

`Ctrl` + `D`

Log out. Corresponds to **exit**. EOF (end of file). Affected by entry in `/etc/profile: ignoreeof=x`. Here, **x** indicates how many times command can be repeated until executed.

`Shift` ↑ + `Ctrl`

Compose key (“Multi_key”).

Bibliography

- [Alm94] ALMESBERGER, Werner: *LILO User's guide*, 1994. – (see file /usr/doc/lilo/user.dvi)
- [Bai97] BAILEY, Edward C.: *Maximum RPM*. Red Hat, 1997. – (ISBN 1-888172-78-9)
- [CAR93] COSTALES, Bryan; ALLMAN, Eric ; RICKERT, Neil: *sendmail*. O'Reilly & Associates, Inc., 1993. – (ISBN 1-56592-056-2)
- [CB96] CHESWICK, William R.; BELLOVIN, Steven M.: *Firewalls und Sicherheit im Internet*. Addison Wesley GmbH, 1996. – (ISBN 3-89319-875-x)
- [CR91] CAMERON, Debra; ROSENBLATT, Bill: *Learning GNU Emacs*. O'Reilly & Associates, Inc., 1991. – (ISBN 0 937175-84-6)
- [CZ96] CHAPMAN; ZWICKY: *Einrichten von Internet Firewalls. Sicherheit im Internet gewährleisten..* O'Reilly & Associates, Inc., 1996. – (ISBN 3-930673312)
- [Daw95] DAWSON, Terry: *Linux NET-2/NET-3 HOWTO*, v2.8, 07 Jan 1995. – (see file /usr/doc/howto/NET-2-HOWTO)
- [FCR93] FANG, Chin; CROSSON, Bob ; RAYMOND, Eric S.: *The Hitchhiker's Guide to X386/XFree86 Video Timing (or, Tweaking your Monitor for Fun and Profit)*, 1993. – (see file /usr/X11/lib/X11/doc/VideoModes.doc)
- [Fri93] FRISCH, Aileen: *Essential System Administration*. O'Reilly & Associates, Inc., 1993. – (ISBN 0-937175-80-3)
- [Gil92] GILLY, Daniel: *UNIX in a nutshell: System V Edition*. O'Reilly & Associates, Inc., 1992. – (ISBN 1-56592-001-5)
- [GMS93] GOOSSENS, Michel; MITTELBACH, Frank ; SAMARIN, Alexander: *The L^AT_EX Companion*. Addison Wesley GmbH, 1993. – (ISBN 3-54199-8)
- [Gri94] GRIEGER, W.: *Wer hat Angst vorm Emacs?*. Addison Wesley GmbH, 1994. – (ISBN 3-89319-620-X)
- [GS93] GARFINKEL, Simson; SPAFFORD, Gene: *Practical UNIX Security*. O'Reilly & Associates, Inc., 1993. – (ISBN 0-937175-72-2)
- [Her92] HEROLD, H.: *UNIX Grundlagen*. Addison Wesley GmbH, 1992. – (ISBN 3-89319-542-8)
- [HHMK96] HETZE, Sebastian; HOHNDEL, Dirk; MÜLLER, Martin ; KIRCH, Olaf: *Linux Anwenderhandbuch*. 6. LunetIX Softfair, 1996. – (ISBN 3-929764-05-9)
- [Hof97] HOFFMANN, Erwin: EMail-Gateway mit qmail. In: *iX* 12 (1997), S. 108ff

- [Hun95] HUNT, Craig: *TCP/IP Netzwerk Administration*. O'Reilly & Associates, Inc., 1995. – (ISBN 3-930673-02-9)
- [Kie95] KIENLE, Micheal: TIS: Toolkit für anwendungsorientierte Firewall-Systeme. In: *iX* 8 (1995), S. 140ff
- [Kir95] KIRCH, Olaf: *LINUX Network Administrator's Guide*. O'Reilly & Associates, Inc., 1995. – (ISBN 1-56592-087-2)
- [Kof95] KOFLER, M.: *Linux*. Addison Wesley GmbH, 1995. – (ISBN 3-89319-796-6)
- [Kop94] KOPKA, Helmut: *LT_EX-Einführung*. Addison Wesley GmbH, 1994. – (ISBN 3-89319-664-1)
- [Kun95] KUNITZ, Ulrich: Sicherheit fast kostenlos: Einrichtung eines kostenlosen Firewall-Systems. In: *iX* 9 (1995), S. 176ff
- [Lam90] LAMB, Linda: *Learning the vi Editor*. O'Reilly & Associates, Inc., 1990. – (ISBN 0-937175-67-6)
- [Lam94] LAMPORT, Leslie: *LT_EX User's Guide and Reference Manual*. Addison Wesley GmbH, 1994. – (ISBN 0-201-52983-1)
- [Lef96a] LEFFLER, Sam: *HylaFAX Home Page*, 1996
- [Lef96b] LEFFLER, Sam: *TIFF Software*, 1996
- [OT92] O'REILLY, Tim; TODINO, Grace: *Manging UUCP and Usenet*. O'Reilly & Associates, Inc., 1992. – (ISBN 0-937175-93-5)
- [Per94] PERLMAN, G.: *Unix For Software Developers*. Prentice-Hall, 1994. – (ISBN 13-932997-8)
- [Pug94] PUGH, K.: *UNIX For The MS-DOS User*. Prentice-Hall, 1994. – (ISBN 13-146077-3)
- [SB92] SCHOONOVER, M.; BOWIE, J.: *GNU Emacs*. Addison Wesley GmbH, 1992. – (ISBN 0-201-56345-2)
- [Sch98] SCHEIDERER, Jürgen: Sicherheit Kostenlos - Firewall mit Linux. In: *iX* 12 (1998)
- [Sto98] STOLL, Clifford: *Kuckucksei; Die Jagd auf die deutschen Hacker, die das Pentagon knackten*. Fischer-TB.-Vlg., 1998. – (ISBN 3596139848)
- [The96] THE XFREE86™-TEAM: *XF86Config(4/5) - Configuration File for Xfree86™*, 1996. – Manual-Page zu XFree86™
- [TSP93] TODINO, Grace; STRANG, John ; PEEK, Jerry: *Learning the UNIX operating system*. O'Reilly & Associates, Inc., 1993. – (ISBN 1-56592-060-0)
- [Wel94] WELSH, Matt: *Linux Installation and Getting Started*. 2. S.u.S.E. GmbH, 1994. – (ISBN 3-930419-03-3)
- [WK95] WELSH, Matt; KAUFMAN, Lars: *Running Linux*. O'Reilly & Associates, Inc., 1995. – (ISBN 1-56592-100-3)

Index

symbols

-probeonly 153
 /etc/conf.modules 231
 /etc/init.d/rc 307
 /etc/inittab 305
 /etc/modules.conf 231
 /etc/rc.config 310
 /etc/rc.config 311
 /sbin/SuSEconfig 310
 /var/lib/apsfilter/SETUP ... 208
Applicware 37, 38, 319
reboot 319
shutdown 319
 /etc/resolv.conf 314
 /etc/host.conf 314
 /etc/hosts 313
 kdm 319
 xdm 319
 SaX2 129, 154
 SuSE Linux
 features 285
 YaST 23
 YaST
 during installation 12
 3COM 3c501 254
 3COM 3c503 247, 254
 3COM 3c505 254
 3COM 3c507 254
 3COM 3c509 254
 3COM 3c515 254
 3COM 3c579 254
 3COM 3c590 254
 3COM 3c900 254
 3Com 238

A

a2ps 212
 a2ps 198
 aaa_base 276
 Adabas 319
 Adabas D 3
 Adaptec

AHA-152x/151x/1505 .. 241,
 252
 AHA-1540/1542 253
 AHA-154x 241
 AHA-274x 242
 AHA-284x 242
 AHA-294x 242
 Adaptec 2920 244, 253
 AdvanSys 242
 aic7xxx 242
 AIRCOP 102
 Aladdin Ghostscript 198
 alice 28
 ALICE 27, 28
alien 284
 Alien 284
 allman 278
 ALSA 169, 170
 alsa-devel 278
 alsacnf 169, 170
 AM53/79C974 242
 AMD 53/79C974 242
 AMD Athlon / Intel Pentium III
 233
 Apache 285
 APM 192
 apmd 193
 apmd 193, 194
 aps 198, 208
 apsfilter .. 70, 71, 208, 209, 211,
 212, 215, 218
 YaST and 70
 apsfilters 212
 Arcad 38
 arena 318
 Argus 319
 Arkeia 319
 AT1700 254
 atd 317
 autofs 316
 autofs 316
 Automounter 316
 awk 1

Aztech CD-ROM 259

B

backups
 creating with YaST 80
 binutils 229
 BIOS 99
 LILO and 103
 book
 using 1
 boot disk
 creating with Unix 33
 boot disks
 creating 112
 creating in DOS 32
 creating with rawrite 33
 creating with setup 32
 installing with 31
 boot.local 309
 boot.sys 101
 booting 23, 99, 323
 BIOS 99
 boot disk 24
 boot graphic 109
 boot loaders 289
 boot manager 101
 boot managers 99
 boot menus 123
 boot sectors 100
 CD2 31
 chaining 101
 CMOS 99
 concepts 100
 cylinder limit 119
 DOS 100
 floppy, from 100
 graphic 25
 disabling 25
 initial ramdisk 288
 inside Windows 125
 kernels after 2.0 121
 LILO 24, 101
 loadlin 24, 122

- logical partitions and 104
 - message file 109
 - methods 24
 - root partition 240
 - sample configurations 114
 - securing 106, 109
 - Windows 100
 - YaST and 71
 - Booting 305
 - bttv 275
 - bus mouse 148
 - BusLogic 243
 - Busmouse
 - Logitech 249
 - buying a PC 296
- C**
- C 229
 - Cabletron E21xx 255
 - Cabletron E21XX 248
 - CardBus *see* hardware, CardBus
 - CardBus network cards based on
 - SMSC 83c170 256
 - cardctl 191
 - Cardmanager 184, 189, 190
 - cardmgr 184
 - cardmgr** 184
 - CD-ROM
 - ATAPI 244
 - ATAPI drives 34
 - Aztech 249, 259
 - EIDE 244
 - Goldstar 249, 260
 - Mitsumi 249, 260
 - Mitsumi FX-001(D) 249, 260
 - Mitsumi Multisession 249, 260
 - Mozart 249, 260
 - Optics Storage 250, 261
 - Panasonic 251, 261
 - parallel port drives 35
 - Philips CM206 250, 261
 - Sanyo 250, 261
 - Sony CDU31A 250, 261
 - Sony CDU33A 250, 261
 - Sony CDU535 251, 261
 - SoundBlaster Pro 16 261
 - supported drives 34
 - CD-ROM drive
 - parallel port . . . *see* hardware,
 - parallel port
 - cds
 - mounting 48
 - check 323
 - Chris Mason 53
 - chroot Environment 61
 - clock chip 150
 - CMOS 99
 - CMOS setup 99
 - color depth 157
 - command not found 321
 - Compose *see* keyboard,
 - Compose key
 - conf.modules 231
 - configuration
 - installation 19
 - net time 318
 - X11 132, 147
 - Configuration
 - Changing 310
 - configuration file 311
 - configuring
 - YaST2 85
 - console 319
 - consoles
 - switching 48, 216, 301
 - core files 287
 - CPU *see* SMP, Number of CPUs
 - crash 323
 - cron 276, 286, 286, 319, 320
 - cron 276
 - cron.daily 320
 - CUPS 97
 - curses 286
 - Cyberscheduler Software 38
 - Cygnus Source-Navigator 38
 - Cyrix . *see* processor, Cyrix 686
- D**
- D-Link DE620 259
 - daemon
 - named 316
 - nntpd 317
 - nscd 316
 - printing 317
 - routed 316
 - rwhod 316
 - DB2 319
 - DCF77 318
 - DE203 255
 - DE204 255
 - DE205 255
 - DEC 129
 - defrag 30, 31
 - dependencies
 - resolving in YaST 65
 - depmod 230
 - dhclient 75
 - DHCP 317
 - diald 316
 - Digital 255
 - Digital DEPCA 255
 - Dirk Hohndel 129
 - disk drives 248
 - docbktls 278
 - DocBook . *see* SGML, DocBook
 - docbook-dsssl-stylesheets
 - 277
 - docbook_3 277
 - docbook_4 277
 - documentation
 - documentation server 317
 - domain name
 - setting 75
 - Duesentrieb 3
 - dummy device 314
 - dump 58
 - dumpe2fs 300
 - Dynamic IP Addresses 314
- E**
- e2fsck 298–301
 - man page 323
 - efscck** 323
 - Eagle 38
 - Eberhard Moenkeberg 3
 - EIDE chipsets 245
 - emacs 39
 - emacs** 3
 - emacs 277
 - Emacs 1, 277
 - emacs-auctex 277
 - emacs-el 277
 - emacs-info 277
 - emacs-nox 277
 - emacs-x11 277
 - emm386.exe 24
 - EtherTeam 16i/32 257
 - ext2 52
- F**
- Fax
 - Hylafax 318
 - faxsetup 318
 - fdisk 29, 30, 56, 103, 110, 111, 116, 117, 298–300
 - fhs 285
 - file permissions 320

- file systems 58
 - File System Hierarchy
 - Standard 285
 - ext2 52
 - FSH 285
 - msdos 54
 - ReiserFS 52
 - repairing 300
 - TeX and 285
 - types 52
 - UMSDOS 54
 - vfat 54
- find 299
- finger 276
- fips 15, 29, 30
- fips.exe 30
- first-time installation
 - fixing the root password ... 20
 - installing packages 17, 22
 - installing software 17
 - modem 21
 - network configuration 20
 - requirements 7
 - starting point 7
- first-time installation of Linux 7
- floppy 248
- floppy disk drives *see* hardware, floppy disk drives
- fonts
 - console 81
- formatting
 - partitions *see* partitions, formatting
- Fortify 38
- free software 1
- Freecom 35
- fsck 58
- fstab 55, 58
- ftape 181
- ftp 75
- FTP
 - server 285
- ftplib 285
- Fujitsu FMV-181/182/183/184 .. 257
- fun 5
- function keys
 - emulating 43
- Future Domain ... 243, 244, 253
- fvwm 138
- G**
 - gamix 169
 - gcc 229
 - GDM 76
 - configuring 76
 - GDT RAID5 controller . *see* ICP Vortex
 - Ghostscript 70, 212, 219
 - print drivers 199
 - gimp-devel 278
 - glibc-devel 229
 - gnadmin 275
 - gnlibs 275
 - GNOME 38
 - GNU 1, 287
 - GNU C/C++ compiler 1
 - GNU Emacs 3
 - GNU Ghostscript 198
 - gnuserv 277
 - gnutils 275
 - Goldstar CD-ROM 249, 260
 - GPL 327, 327
 - gpm 82, 139, 313
 - GPM 313
 - graphics cards
 - FireGL 8
 - grep 299
 - groups
 - managing with YaST 79
 - gs** 212
 - gs_fonto 198
 - gs_fonts 199
 - gs_lib 199
 - gs_x11 198
 - gv 219
 - gxmhtml 275
 - gzip 41
- H**
 - halt 309
 - Hans Reiser 53
 - hard disk
 - parameters 245
 - partitioning 50
 - hardware
 - adding 70
 - AGP 165
 - CardBus 183
 - cards 165
 - CD-ROM drives 34
 - external modems 178
 - floppy disk drives 176
 - integration 165
 - internal modems 178
 - IrDA *see* IrDA
 - ISA 165
 - ISA PnP *see* ISA PnP
 - laptop *see* notebook
 - LS120 drives 176
 - modems 178
 - modules 9
 - notebook *see* notebook
 - parallel port 171
 - PC cards 183
 - PCI 165
 - plugnplay 166
 - streamers . *see* hardware, tape drives
 - tape drives 180
 - Winmodem 179
 - YaST2 and 86
 - ZIP drives 177
- Hardware
 - Floppy disk drive 248
 - Plug and Play 308
 - USB *see* USB
- hdparm 194
- help 317
 - info pages 287
 - man pages 287
 - YaST 44
- hibernation 192
- HOME 271
- horizontal frequency 149
- host name
 - setting 75
- howto 195
- HP 10/100 VG-AnyLAN ... 257
- HP PCLAN 257
- HP PCLAN+ 257
- htdig 38
- http-rman 317
- httpd 317
- I**
 - I18N 302
 - i41 76, 310
 - i82365 189
 - i82557/i82558 256
 - IBM Thinkpad 248
 - ICL EtherTeam 257
 - ICP Vortex controller
 - errors in installation 12
 - idedma 277
 - identifier 157
 - ifconfig 299, 314
 - ifport 190

- important keys 335
- imwheel 162
- inetd 20, 75, 315
- info** 287
- Info 287
- info pages 287
- Infrared Data Association ... *see* IrDA
- infrared support *see* IrDA
- init 2, 299, 305, 305–308
 - scripts 307
- initrd 289, 313
 - boot loaders and 289
 - creating 291
 - linuxrc and 289
 - SCSI driver 231
- inittab 305
- INN 317
- inodes
 - density 54
 - YaST configuration of 54
- insmod 230, 231, 237
 - parameters 237
- insserv 308
- installation
 - adjusting in YaST 45
 - base configuration 19
 - boot methods 24
 - configuration selection 16
 - first-time installation 7
 - logging in 22
 - login 22
 - mouse 21
 - PCMCIA 190
 - selecting a kernel 18
 - software selection 16
 - starting YaST 12
- Installation 7
- installing
 - ALICE 27
 - automated 27
 - boot disk 31
 - FTP 26
 - hardware modules 9
 - linuxrc 8
 - net as source medium 26
 - NFS 26
 - problems 8, 11
 - selecting keyboard 9
 - selecting language 9
 - special installations 26
 - starting 11
 - welcome screen 7
 - YaST 7, 67
 - YaST1 7
- Insure++ 38
- Intel EtherExpress 16 256
- Intel EtherExpressPro 256
- Intel EtherExpressPro 100 .. 256
- Internet Services Daemon
 - xinetd 315
- io controlled clocks 318
- IP Addresses
 - Dynamic 314
- IP forwarding 314
- iputils 276
- IRC 315
- irda 195
- IrDA 194
 - printer 195
- ISA cards 166
- ISA PnP
 - initialize 313
- isapnp 166–168
 - isapnp 166, 168
- ISAPnP 166
- iso-8859 159
- ISP16 250, 260
- J**
- jade_dsl 276
- K**
- kardinfo 191
- kbatmon 194
- KDE 38, 284
- kdelibs-devel 278
- KDM 23, 76
 - configuring 76
- kernel 229
 - compiling 229
 - configuration 232
 - modules 230
 - parameters 237
 - printer driver 269
 - select SMP kernel 19
- Kernel
 - Debugging 312
 - Number of CPUs 240
 - Serial console 240
 - Sysrq 312
- kernel daemon ... 231, 308, 313
- kernel panic 19
- kernel too large 233
- kernel-source 229
- kernel 231, 313
- kernels
 - selecting 96
 - YaST and 71
- keyboard
 - CapsLock 312
 - Compose key 335
 - configuring in YaST 45
 - delay 312
 - layout 301
 - map 312
 - NumLock 312
 - repeat 312
 - selecting for install 9
- kfontmanager 161
- kikbd 163
- kmix 169
- kmod 200
- Kmod 172, 231
- knfsd 53
- knfsd 316
- kpanel 194
- krpm 284
- kterm 81
- L**
- L10N 302
- Lance 247
- LANG 312
- language
 - selecting for install 9
 - selecting in YaST 45
- Language
 - Configuration 312
- laptop 183
- LC_* 312
- LDAP server 317
- less 299, 321
- libc 283
- libc5 298
- License 327
- LILO 18, 19, 24–26, 31, 37, 53, 72–74, 100–122, 132, 231, 234, 235, 237, 239, 252, 253, 274, 288, 289, 291, 295, 306
 - BIOS and 103
 - boot disks 112, 113
 - boot manager, with another ... 104
 - boot sector, in 103
 - components 102
 - configuring 24, 104

- configuring in YaST 72
- DOS/boot Win 95/98 114
- floppy, from 103
- installing 110
- lilo.conf 105
- locations for 103
- map file 102
- MBR 104
- message file 102, 109
- options 107
- other systems 107
- overview 101
- parameters 106, 237
- passwords 109
- problems with 117
- reinstalling 103, 109, 110
- sample configurations ... 114
- troubleshooting 118
- uninstalling 111
- YaST and 104
- Linus Torvalds 3, 5
- Linux
 - distributions 1
 - kernel 1
 - starting inside Windows .. 125
 - what is 1
- Linux Development Platform
 - Specification 286
- Linux Standard Base 285
- linuxrc 4, 7–13, 27, 32, 35, 190, 273, 289, 290, 292, 294–297
- initrd and 289
- loadlin 24, 25, 36, 38, 101, 119, 121–124, 127, 167, 239, 274, 288, 289, 294
- boot menus 123
- configuring 24
- problems 36
- Loadlin doesn't start 36
- loadlin.exe 3, 290
- local 312
- locale 312
- locate 286, 320
- log files 321
- Logical Volume Manager
 - configuring in YaST 60
 - YaST2 89
- login 319
- configuring 76
- Login
 - remote 321
- Logitech 148
- Logitech Busmouse 249
- logview 275
- loopback 313
- lpc 206
- lpd 204, 218, 317
- LPD 97
- lpq 205
- lpr 207, 208
- lprm 205
- lprold 97, 199, 218
- lprsetup 208
- lprsetup 70, 208
- LS120 drives *see* hardware, LS120 drives
- LSB .. *see* Linux Standard Base
- lsmod 231
- lspci 157, 158
- LUN 241
- lynx 318
- M**
- MAD16 250, 260
- mail server
 - Postfix 315
- make menuconfig 232
- make zImage 111
- man pages 287
- Masquerading
 - IP Forwarding 314
- Master Boot Record 99
- LILO 104
- MBR
 - see* Master Boot Record ... 99
- mc 284
- mc 283
- memory
 - not recognized 240
 - protect 239
 - reserve 239
- memory check 239
- MemTest86 105
- Midnight Commander 283
- MIT 129
- Mitsumi CD-ROM 249, 260
- Mitsumi FX-001(D) .. 249, 260
- mkfontdir 161
- mkfs 299
- mkswap 299
- mod_php 277
- mod_php4 277
- modeline 158
- modem 313
- modprobe ... 230, 231, 237, 252
- parameters 237
- modules 230
- compile 234
- handling 230
- loading 294
- parameters 294
- modules.conf 231
- monitors 149
- mount 176, 177, 299
- mount points 53, 58
- mounting
 - mount points 52
- mouse 313
- Bus 147
- HiTablet 147
- Logitech 147
- Logitech (MouseMan) ... 147
- Microsoft 147
- MM series 147
- Mouse Systems 147
- PS/2 147
- mouse buttons 148
- mouse type 148
- Mozart 250, 260
- Mozart CD-ROM 249, 260
- Mozilla 38
- mt 180
- multiprocessor systems *see* SMP
- Multisession CD-ROM 249, 260
- MySQL 278, 319
- mysql-Max 278
- N**
- Name Service Caching Daemon
 - Initialize 316
- ncpfs 76
- NCR 5380 244, 253
- NCR 53c400 244
- NCR 53C400 253
- NCR 53c406a 244
- NCR 53C810 238
- ncurses 286
- ncurses 286
- NE1000/2000 258
- NE2000 251
- Netscape 37–39
- netstat 299
- network
 - configuring with YaST 74
 - dummy device 314
 - YaST2 and 86
- network card
 - AM7990 chipset 258

- network cards 247
 - 3COM 3c501 254
 - 3COM 3c503 254
 - 3COM 3c505 254
 - 3COM 3c507 254
 - 3COM 3c509 254
 - 3COM 3c515 254
 - 3COM 3c579 254
 - 3COM 3c590 254
 - 3COM 3c900 254
 - AT1700 254
 - Cabletron 255
 - D-Link DE620 259
 - DE10x 255
 - DE20 255
 - DE203 255
 - DE204 255
 - DE205 255
 - DE42 255
 - DE425 255
 - DE434 255
 - DE435 255
 - DE450 255
 - DE500 255
 - DEC EtherWORKS 255
 - Digital 255
 - Digital DEPCA 255
 - E21xx 255
 - EtherBlaster 258
 - EtherTeam 16i/32 257
 - EtherWORKS 3 255
 - Fujitsu FMV-181/182/183/184
257
 - HP 10/100 VG-AnyLAN . 257
 - HP 27245 257
 - HP 27247B 257
 - HP 27252A 257
 - HP 27xxx 257
 - HP PCLAN 257
 - HP PCLAN+ 257
 - IBM Token Ring 259
 - ICL EtherTeam 257
 - Intel EtherExpress 16 256
 - Intel EtherExpress Pro 100 ...
256
 - Intel EtherExpressPro ... 256
 - Lance 258
 - Novell NE1000/2000 258
 - SMC 9194 258
 - SMC 9432TX EtherPower II .
256
 - SMC Ultra 258
 - Token Ring 259
 - WD80x3 258
 - Western Digital 258
 - network monitor
 - Argus 319
 - NFS
 - group IDs 316
 - Kernel NFS daemon 316
 - ReiserFS 53
 - server 315
 - user IDs 316
 - NI6510 258
 - NIS 275
 - nkit 283
 - nkitb 276
 - NNTP 317
 - NNTP server 315
 - nntpd 317
 - nobody 320
 - notebook 183
 - IrDA *see* IrDA
 - PCMCIA 313
 - power management . *see* APM
 - Novell NE1000/2000 258
 - Novell Server emulation ... 318
- O**
- Omnibook 248
 - Optics Storage CD-ROM .. 250,
261
 - Oracle 8.1.7 38
 - OSS 169
- P**
- package
 - a2ps 198
 - aaa_base 276
 - alice 28
 - allman 278
 - alsa-devel 278
 - apmd 193, 194
 - aps 198, 208
 - autofs 316
 - binutils 229
 - bttv 275
 - cron 276
 - dhclient 75
 - docbktls 278
 - docbook-dsssl-stylesheets
277
 - docbook_3 277
 - docbook_4 277
 - emacs 277
 - emacs-auctex 277
 - emacs-el 277
 - emacs-info 277
 - emacs-nox 277
 - emacs-x11 277
 - fhs 285
 - finger 276
 - ftplib 285
 - gcc 229
 - gimp-devel 278
 - glibc-devel 229
 - gnadmin 275
 - gnlibs 275
 - gnuserv 277
 - gnutils 275
 - gs_fonto 198
 - gs_fonts 199
 - gs_lib 199
 - gs_x11 198
 - gxhtml 275
 - howto 195
 - i4l 76, 310
 - idedma 277
 - imwheel 162
 - iputils 276
 - irda 195
 - isapnp 166, 168
 - jade_dsl 276
 - kdelibs-devel 278
 - kernel-source 229
 - knfsd 316
 - libc 283
 - lprold 97, 199, 218
 - mod_php 277
 - mod_php4 277
 - mysql-Max 278
 - ncpfs 76
 - ncurses 286
 - nkit 283
 - nkitb 276
 - pcmcia 190, 191
 - PCMCIA 184
 - pg_datab 276
 - plp 203
 - postgres 272, 276
 - psgml 277
 - reiserfs 53
 - rpm 283
 - rsh 276
 - rwho 276
 - rzsz 276
 - samba 76

- sax 129
- sax2 129
- sdb_en 275
- sendmail 310
- syslinux 31
- talk 276
- tcl 233
- telnet 276
- tk 233
- ttmkfdir 160
- uucp 287
- wget 282
- xdevel 286
- xdevel33 286
- xf86 233
- xfsetup 129
- xkeycaps 163
- xntp 318, 319
- xvga16 129
- yast2-config-tune ... 277
- ypclient 75, 276
- ypserv 275
- yudit 160
- package selection 16
- packages
 - compiling 282
 - installing 279
 - RPMs 278
 - uninstalling 279
 - updating 274
 - verifying 279
- PAM 79
- Panasonic CD-ROM .. 251, 261
- parallel port
 - architecture-specific 269
 - ATAPI CD-ROM 270
 - ATAPI floppy disks 270
 - ATAPI streamer drive 270
 - generic ATAPI device 270
 - IDE devices 269
 - IDE hard disk 270
 - IDE protocol driver 269
 - kernel parameter ... 251, 269
- parallel port seehardware,
 - parallel port 171
- Paride 269
- parport subsystem 171
- partitions
 - boot 53
 - configuring 56
 - creating 50
 - DOS 54
 - expert partitioning 39
 - extended 51
 - FAT32 30
 - file systems 52
 - formatting 55
 - logical 51
 - mounting 52
 - partition table 99
 - primary 51
 - problems partitioning 31
 - reducing size of DOS
 - partitions 29
 - reducing Windows partitions .
 - 30
 - swap 39, 50
 - tips and tricks 28, 36
 - types 36
 - UMSDOS 54
 - Windows 54
 - YaST2 and 88
- passwords 79
- PATH 3
- Patrick Volkerding 5
- PC cards *see* Hardware, PC cards
- PCI cards 165
- pcmcia 190, 191
- PCMCIA 183, 313
 - cardmgr 184
 - configuration 184
 - Ethernet 184
 - help programs 191
 - IDE 185
 - installation 190
 - ISDN 185
 - modem 185
 - problems 188
 - scheme 186
 - SCSI 185
 - software 184
 - Token Ring 184
- PCMCIA 184
- pep 38
- perl 1, 320
- permissions 320
- pg_datab 276
- Philips CM206 250, 261
- PLIP
 - parallel port ... *see* hardware, parallel port
- plp 203
- plugáplay 166
- pnpdump 166–168
- portmap 20, 75
- portmapper 275, 315
- ports
 - parallel 199
 - serial 203
 - USB 202
- postfix 315
- Postfix 315
- postgres 272, 276
- PostgreSQL 272, 276
- power management 192
- printer
 - kernel parameter 251
 - parallel port ... *see* hardware, parallel port
- Printer
 - Connecting to a Windows
 - Printserver 76
- PRINTER 207
- printer driver 269
- printers
 - configuring in YaST 70
 - connecting 199
 - GDI 197
 - GDI printers 197
 - languages 197
 - network 205
 - configuring 216
 - serial 203
 - supported models 198
 - USB 202
 - Windows only 197
- printing 197
 - lpd 317
 - apsfilter 70, 211
 - configuring 208
 - controlling queues 206
 - creating filters 220
 - CUPS 97
 - drivers 199
 - languages 197
 - listing jobs 205
 - lpc 206
 - lpq 205
 - lpr 207
 - lprm 205
 - lprsetup 208
 - managing 204, 205
 - network 215
 - configuring 216
 - filtering 215

- prefilter 217
- testing 218
 - network printer 215
 - network printers 205
 - packages required 198
 - print filter 197
 - print queue 197
 - print server 205, 215
 - print spooler 197
 - printserver unit 215
 - problems 205
 - removing jobs 205
 - requirements 197
 - security 205
 - starting 207
 - testing 210
 - tools 205
- Printing
 - Samba 76
- Pro Audio Spectrum .. 241, 250
- Pro Audio Spectrum 16 250
- processor
 - Cyrix 686 19
- programming
 - core files 287
- protect I/O address 239
- Proxy
 - FTP 318
 - Gopher 318
 - HTTP 318
- psgml 277
- R**
- radius 317
- RAID 94
 - soft 94
- RAM *see* memory
- RAMDAC 150
- ramdisk
 - initial ramdisk 313
- rawrite 33
- rawrite.exe 33
- rc 307
- rc.config 82, 311
 - editing with Yast2 87
 - /etc/rc.config 310
 - rc.config 311
 - rcp.ugidd 316
- reboot 23, 239, 319
- reiserfs 53
- ReiserFS 52
 - NFS and 53
 - RAID and 53
- reiserfsck 53
- rescue disk 296
- rescue disks
 - creating in YaST 72
- rescue system 296
- reset 239
- resolv.conf 288
- rlogin 315, 321
- rmail 315
- rmmmod 230
- root 20
- Root
 - Einloggen, remote 321
- root partition 240
- route 299
- Router
 - IP Forwarding 314
- rpc.mountd 315
- rpc.nfsd 75, 315
- rpm 282
- rpm** 278
- rpm 283
- RPM 278
 - database 278, 320
- rebuilding 282
 - dependencies 279
 - installing 279
 - queries 280
 - SRPMS 283
 - tools 283
 - uninstalling 279, 280
 - updating 279
 - verify 281
 - verifying 279
- rsh 276
- runlevel 305
 - changing 307
- runlevels
 - editor in YaST2 87
- rwho 276
- rwhod 316
- rzsz 276
- S**
- samba 76
- Samba 277, 318
- Sanyo CD-ROM 250, 261
- sax 129
- SaX 81, 129–132, 134, 138–147
- sax2** 129
- sax2 129
- SaX2 129, 130, 132–134, 136–138, 154, 155, 157, 158, 161, 162
- scheme .. *see* PCMCIA, scheme
- screen resolution 157
- SCSI
 - Adaptec
 - AHA-152x/151x/1505 . 252
 - Adaptec AHA-1540/1542 253
 - AdvanSys 242
 - AHA-152x/151x/1505 ... 241
 - AHA-154x 241
 - AHA-274x 242
 - AHA-284x 242
 - AHA-294x 242
 - AM53/79C974 242
 - Future Domain . 243, 244, 253
 - LUN 241
 - NCR 5380 244, 253
 - NCR 53c400 244
 - NCR 53C400 253
 - NCR 53c406a 244
 - Seagate ST01/02 244
 - streamer 241
 - TMC-16x0 243, 253
 - TMC-885/950 244
 - Trantor T128/128F/228 .. 244
 - Trantor T130B 253
- sdb_en 275
- Seagate ST01/02 244
- secure shell daemon 318
- sed 1
- sendmail 75
- sendmail 310
- Sendmail 287
- sendmail -q 315
- Serial console 240
- series
 - a 53, 64
 - al 65, 184
 - ap ... 162, 166, 193, 198, 199
 - D 229
 - doc 275, 285
 - n .. 75, 76, 195, 199, 203, 316
 - sgm 277
 - x 129
 - xap 160, 163
 - xsrv 129
 - zq 282, 283
- server.exe 101
- set up time 318
- setserial 171, 179

- setup 32, 122, 123
- Setup 32
- SETUP 70, 208
- setup.exe 4
- SGML
 - DocBook 278
- shutdown 299, 319
- Shutdown 23
- SMC 9194 258
- SMC 9432TX EtherPower II 256
- SMC EPIC/100
 - 83C170/83c171/83c172/83c175
256
- SMC Ultra 258
- SMP
 - Number of CPUs 240
 - select kernel 19
- SMTP 315
- SNiFF+ 38
- Soft RAID *see* YaST2,Soft
RAID
- software
 - compiling 278
 - software selection 16
- Sony CDU31A 250, 261
- Sony CDU33A 250, 261
- Sony CDU535 CD-ROM .. 251,
261
- sound
 - AD1816 Chip 262
 - Aztech Sound Galaxy 267
 - configuring 169
 - Ensoniq SoundScape 267
 - Gravis Ultrasound 263
 - MediaTrix AudioTrix Pro 267
 - MPU401 264
 - OPL3 265
 - OPL3-SA1 265
 - OPL3-SAx 265
 - Personal Sound System
(ECHO ESC614) 266
 - Pro Audio Spectrum 266
 - S3 Sonic Vibes 267
 - Sound Blaster and clones 266
 - Sound Blaster DSP chipsets ..
268
 - Turtle Beach Maui and Tropez
264
 - Turtle Beach Maui, Tropez,
Tropez Plus 268
 - Turtle Beach MultiSound 264
 - Turtle Beach Pinnacle/Fiji 265
 - UART401 268
 - UART6850 268
 - YMF71x 265
 - Sound
 - AD1848/CS4248 Chip (MSS)
262
 - Creative Ensoniq 1371 chipset
263
 - Crystal 423x chipsets 263
 - Ensoniq 1370 chipset 263
 - Generischer OPLx driver . 262
 - MAD16 264
 - Turtle Beach
 - Classic/Monterey/Tahiti 264
 - sound card
 - Pro Audio Spectrum 250
 - source
 - compiling 282
 - packages 278
 - source code 64
 - spindown 194
 - spm 282
 - squid 317
 - SRPMS 283
 - ssh 81
 - standby 192
 - StarOffice 38
 - startx 76, 138
 - streamer
 - SCSI 241
 - streamers *see* hardware,
streamers
 - SuSE 285
 - SuSEconfig ... 20, 82, 186, 276,
298, 302, 303, 310, 311,
313, 314, 316, 318, 321
 - SuSEconfig 310, 311
 - SuSEConfig 310
 - susehelp 38
 - SuSE Linux 285
 - susewm 78
 - YaST 77
 - suspend 192
 - Sybase 38
 - Syn Flood Protection 314
 - syslinux 289, 290
 - syslinux 31
 - SYSLINUX 31, 32, 239
 - syslog 299
 - Sysrq *see* Kernel, Sysrq
system
 - information 293
 - limiting resource use 287
 - making local adjustments 302
 - rescue
 - starting 298
 - rescuing 296
 - updating 271
 - with YaST 272
 - updating with YaST 69
 - system administration
 - YaST 69
 - system configuration 311
 - system is too large 233

T

 - talk 276
 - tape drives .. *see* hardware, tape
drives
 - tar 41, 287
 - tc1 233
 - telnet 75, 299, 315, 321
 - telnet 276
 - temporary files
 - delete 321
 - deleting at bootup 321
 - termcap 286
 - testing a PC 296
 - teTeX 285
 - text console 319
 - The XFree86 Project 129
 - time zones 311, 318
 - setting with YaST 81
 - tk 233
 - TMC-16x0 243, 253
 - TMC-885/950 244
 - Token Ring 259
 - Trantor T128/128F/228 244
 - Trantor T130B 253
 - ttmkfdir 160
 - ttmkfdir 160

U

 - ulimit 287
 - options 287
 - Ultrastor 238
 - unfsd 53
 - Unicode 160
 - Universal Serial Bus .. *see* USB
 - update
 - patch CD 87
 - updating
 - packages 274
 - preparing 271
 - system 271

- with YaST 272
- USB 173
 - Example configuration of mouse and keyboard .. 176
 - Moduls 173
- useradd 79
- userdel 79
- users
 - managing with YaST 78
- uucp 287
- uucp 287
- V**
- Variable
 - HOME 271
 - LANG 312
 - LC_* 312
 - PATH 3
 - PRINTER 207
- variables
 - environment 302
- vertical frequency 149
- VG-AnyLAN 257
- vi 1, 299
- virtual screen 157
- virtuoso 38
- virus 25, 102
- W**
- Wabi 38
- WD80x3 258
- web servers
 - setting up 285
- Western Digital WD80x3 .. 258
- wget 318
- wget 282
- Windows
 - Samba 318
- Windows Explorer 125
- wvdial 185, 196
- wvdial 185
- X**
- X 153
- X -probeonly 152, 153
- X Consortium, Inc. 1, 129
- X Window System .. 1, 129, *see* X11, 154
- X11 129
 - character fonts 159
 - configuration 132, 147
 - keyboard 148
 - mice 147
- monitors 149
- X server 150
 - display manager 319
 - driver 158
 - fonts 159
 - framebuffer 132
 - graphics cards 150
 - logfile 136
 - mkfontdir 159
 - optimization 154
 - shutdown 319
 - start 138
 - troubleshooting 132
 - True type font 159
 - ttmkfdir 159
 - VGA 133
 - VGA16 133
- X11R1 129
- X11R6.4 129
- x-devel 286
- x-devel33 286
- xdm 313
- XDM 76, 271
 - configuring 76
- xf86 233
- XF86_8514 131
- XF86_Mach32 131
- XF86_Mach8 131
- XF86_S3 131
- xf86config . 129, 147, 150, 152, 154
- XF86Config 147
 - clocks 157
 - depth 157
 - device 155, 156
 - device section 157
 - files 155
 - input device 155
 - modeline 155, 157
 - modes 155, 157
 - monitor 155, 156
 - monitor section 158
 - screen 155
 - screen section 156
 - server flags 155
 - server layout 155
 - subsection
 - display 157
- XF86Setup 129, 147
- XFree86 1, 130
 - history 129
- XFree86™ 129
- xfsetup 129
- xinetd 315
- xkeycaps 163
- xkeycaps 163
- xlsfonts 161
- xntp 318, 319
- xrpm 284
- xset -fp 159
- xset +fp 159
- xsetpointer 162
- xterm 81, 303
- xvga16 129
- xvidtune 135
- Y**
- yast** 23, 43
- YaST .. 2, 4, 7, 8, 12–24, 31, 37, 43–46, 48, 51, 55–74, 78–83, 104, 110, 114, 119, 129, 130, 138, 142, 184, 186, 191, 195, 198, 229, 272–274, 276, 282–284, 290, 297, 302, 310–314
 - adjusting installation 45
 - backups 80
 - booting 71
 - checking package selections .. 67
 - configuration file 82
 - configurations 63
 - console fonts 81
 - deleting packages 69
 - dependencies 65
 - file system error behavior . 55
 - file systems 52, 58
 - formatting partitions 55
 - fstab 55
 - gpm 82
 - group administration 79
 - hard disk configuration 56
 - hardware 70
 - help 44
 - inode Density 54
 - installation medium 45
 - installation size 61
 - installing 67
 - installing new system 7
 - installing packages 67
 - kernels 71
 - keyboard configuration 45
 - language selection 45
 - LILO configuration 72
 - loading configuration 62

- login configuration 76
 - LVM 60
 - mount points 52, 53, 58
 - navigating 43
 - network configuration 74
 - partitioning hard disks 50
 - printers 70
 - rescue disks 72
 - saving configuration 62
 - searching packages 67
 - selecting packages 63
 - selecting series 63
 - starting 43
 - susewm 77
 - system administration 69
 - system security 81
 - time zone 81
 - uninstalling packages 69
 - updating system 69
 - user administration 78
 - YaST1 ... 11, 28, 207, 211, 272
 - YaST2 . 7, 12, 25, 60, 70, 85–88, 90–95, 97, 119, 169, 170, 190, 198, 207, 212, 272, 276, 277
 - configuring system 85
 - CUPS 97
 - hardware 86
 - kernel selection 96
 - Logical Volume Manager . 89
 - LVM 89
 - network 86
 - partitioner 88
 - patch CD update 87
 - RC.Config editor 87
 - runlevel editor 87
 - Soft RAID 94
 - yast2-config-tune 277
 - YP
 - server 316
 - ypclient 75, 276
 - ypserv 275
 - yudit 160
 - yudit 160
- Z**
- ZIP drive
 - parallel port ... *see* hardware, parallel port
 - ZIP drives ... *see* hardware, ZIP drives