Debian GNU/Linux Installation Guide

Debian GNU/Linux Installation Guide

Copyright © 2004, 2005, 2006, 2007 the Debian Installer team

This document contains installation instructions for the Debian GNU/Linux 4.0 system (codename “etch”), for the PA-RISC (“hppa”) architecture. It also contains pointers to more information and information on how to make the most of your new Debian system.

Warning

This installation guide is based on an earlier manual written for the old Debian installation system (the “boot-floppies”), and has been updated to document the new Debian installer. However, for hppa, the manual has not been fully updated and fact checked for the new installer. There may remain parts of the manual that are incomplete or outdated or that still document the boot-floppies installer. A newer version of this manual, possibly better documenting this architecture, may be found on the Internet at the `debian-installer` home page (http://www.debian.org/devel/debian-installer/). You may also be able to find additional translations there.

This manual is free software; you may redistribute it and/or modify it under the terms of the GNU General Public License. Please refer to the license in Appendix F.

Table of Contents

Installing Debian GNU/Linux 4.0 For hppa viii
1. Welcome to Debian 1
1.1. What is Debian? 1
1.2. What is GNU/Linux? 2
1.3. What is Debian GNU/Linux? 3
1.4. Getting Debian 3
1.5. Getting the Newest Version of This Document 4
1.6. Organization of This Document 4
1.7. About Copyrights and Software Licenses 4
2. System Requirements 6
2.1. Supported Hardware 6
2.1.1. Supported Architectures 6
2.1.2. CPU, Main Boards, and Video Support 7
2.1.3. Graphics Card Support 7
2.1.4. Multiple Processors 8
2.2. Installation Media 8
2.2.1. CD-ROM/DVD-ROM 8
2.2.2. Hard Disk 8
2.2.3. Network 8
2.2.4. Un*x or GNU system 9
2.2.5. Supported Storage Systems 9
2.3. Peripherals and Other Hardware 9
2.4. Purchasing Hardware Specifically for GNU/Linux 9
2.4.1. Avoid Proprietary or Closed Hardware 9
2.5. Memory and Disk Space Requirements 10
2.6. Network Connectivity Hardware 10
2.6.1. Drivers Requiring Firmware 10
3. Before Installing Debian GNU/Linux 11
3.1. Overview of the Installation Process 11
3.2. Back Up Your Existing Data! 12
3.3. Information You Will Need 12
3.3.1. Documentation 12
3.3.1.1. Installation Manual 12
3.3.1.2. Hardware documentation 12
3.3.2. Finding Sources of Hardware Information 12
3.3.3. Hardware Compatibility 14
3.3.4. Network Settings 14
3.4. Meeting Minimum Hardware Requirements 15
3.5. Pre-Partitioning for Multi-Boot Systems 15
3.6. Pre-Installation Hardware and Operating System Setup 16
4. Obtaining System Installation Media 18
4.1. Official Debian GNU/Linux CD-ROM Sets 18
4.2. Downloading Files from Debian Mirrors 18
4.2.1. Where to Find Installation Images 18
4.3. Preparing Files for TFTP Net Booting 18
4.3.1. Setting up a BOOTP server 19
4.3.2. Setting up a DHCP server 19
4.3.3. Enabling the TFTP Server 20

4.3.4. Move TFTP Images Into Place ... 21
4.4. Automatic Installation ... 21
4.4.1. Automatic Installation Using the Debian Installer ... 21
5. Booting the Installation System ... 22
5.1. Boot Parameters ... 22
5.1.1. Debian Installer Parameters ... 22
5.1.1.1. Using boot parameters to answer questions ... 24
5.1.1.2. Passing parameters to kernel modules ... 25
5.1.1.3. Blacklisting kernel modules ... 26
5.2. Troubleshooting the Installation Process ... 26
5.2.1. CD-ROM Reliability ... 26
5.2.1.1. Common issues ... 26
5.2.1.2. How to investigate and maybe solve issues ... 27
5.2.2. Boot Configuration ... 28
5.2.3. Interpreting the Kernel Startup Messages ... 28
5.2.4. Reporting Installation Problems ... 28
5.2.5. Submitting Installation Reports ... 29
6. Using the Debian Installer ... 31
6.1. How the Installer Works ... 31
6.2. Components Introduction ... 31
6.3. Using Individual Components ... 34
6.3.1. Setting up Debian Installer and Hardware Configuration ... 34
6.3.1.1. Check available memory / low memory mode ... 34
6.3.1.2. Selecting Localization Options ... 35
6.3.1.3. Choosing a Keyboard ... 35
6.3.1.4. Looking for the Debian Installer ISO Image ... 36
6.3.1.5. Configuring Network ... 36
6.3.2. Partitioning and Mount Point Selection ... 37
6.3.2.1. Partitioning Your Disks ... 37
6.3.2.2. Configuring Multidisk Device (Software RAID) ... 39
6.3.2.3. Configuring the Logical Volume Manager (LVM) ... 41
6.3.2.4. Configuring Encrypted Volumes ... 42
6.3.3. Setting up the System ... 45
6.3.3.1. Configuring Your Time Zone ... 46
6.3.3.2. Configuring the Clock ... 46
6.3.3.3. Setting Up Users And Passwords ... 46
6.3.3.3.1. Set the Root Password ... 46
6.3.3.3.2. Create an Ordinary User ... 46
6.3.4. Installing the Base System ... 47
6.3.5. Installing Additional Software ... 47
6.3.5.1. Configuring apt ... 47
6.3.5.2. Selecting and Installing Software ... 47
6.3.6. Making Your System Bootable ... 48
6.3.6.1. Detecting other operating systems ... 49
6.3.6.2. **palo**-installer ... 49
6.3.6.3. Continue Without Boot Loader ... 49
6.3.7. Finishing the Installation ... 49
6.3.7.1. Finish the Installation and Reboot ... 49
6.3.8. Miscellaneous ... 49
6.3.8.1. Saving the installation logs ... 50
6.3.8.2. Using the Shell and Viewing the Logs ... 50

6.3.8.3. Installation Over the Network ... 50
7. Booting Into Your New Debian System ... 52
7.1. The Moment of Truth ... 52
7.2. Mounting encrypted volumes ... 52
7.2.1. dm-crypt ... 52
7.2.2. loop-AES ... 53
7.2.3. Troubleshooting ... 53
7.3. Log In ... 54
8. Next Steps and Where to Go From Here ... 55
8.1. Shutting down the system ... 55
8.2. If You Are New to Unix ... 55
8.3. Orienting Yourself to Debian ... 55
8.3.1. Debian Packaging System ... 55
8.3.2. Application Version Management ... 56
8.3.3. Cron Job Management ... 56
8.4. Further Reading and Information ... 56
8.5. Setting Up Your System To Use E-Mail ... 57
8.5.1. Default E-Mail Configuration ... 57
8.5.2. Sending E-Mails Outside The System ... 57
8.5.3. Configuring the Exim4 Mail Transport Agent ... 58
8.6. Compiling a New Kernel ... 59
8.6.1. Kernel Image Management ... 59
8.7. Recovering a Broken System ... 60
A. Installation Howto ... 62
A.1. Preliminaries ... 62
A.2. Booting the installer ... 62
A.2.1. CDROM ... 62
A.2.2. Booting from network ... 62
A.2.3. Booting from hard disk ... 62
A.3. Installation ... 63
A.4. Send us an installation report ... 63
A.5. And finally... ... 64
B. Automating the installation using preseeding ... 65
B.1. Introduction ... 65
B.1.1. Preseeding methods ... 65
B.1.2. Limitations ... 65
B.2. Using preseeding ... 66
B.2.1. Loading the preconfiguration file ... 66
B.2.2. Using boot parameters to preseed questions ... 67
B.2.3. Auto mode ... 67
B.2.4. Aliases useful with preseeding ... 69
B.2.5. Using a DHCP server to specify preconfiguration files ... 69
B.3. Creating a preconfiguration file ... 70
B.4. Contents of the preconfiguration file ... 71
B.4.1. Localization ... 71
B.4.2. Network configuration ... 71
B.4.3. Mirror settings ... 72
B.4.4. Partitioning ... 73
B.4.5. Partitioning using RAID ... 74
B.4.6. Clock and time zone setup ... 75

B.4.7. Apt setup ... 76
B.4.8. Account setup ... 76
B.4.9. Base system installation ... 77
B.4.10. Boot loader installation ... 77
B.4.11. Package selection ... 78
B.4.12. Finishing up the first stage install ... 78
B.4.13. X configuration ... 79
B.4.14. Preseeding other packages ... 79
B.5. Advanced options ... 79
B.5.1. Running custom commands during the installation ... 80
B.5.2. Using preseeding to change default values ... 80
B.5.3. Chainloading preconfiguration files ... 80
C. Partitioning for Debian ... 82
C.1. Deciding on Debian Partitions and Sizes ... 82
C.2. The Directory Tree ... 82
C.3. Recommended Partitioning Scheme ... 83
C.4. Device Names in Linux ... 84
C.5. Debian Partitioning Programs ... 85
C.5.1. Partitioning for PA-RISC ... 85
D. Random Bits ... 86
D.1. Linux Devices ... 86
D.1.1. Setting Up Your Mouse ... 87
D.2. Disk Space Needed for Tasks ... 87
D.3. Installing Debian GNU/Linux from a Unix/Linux System ... 88
D.3.1. Getting Started ... 88
D.3.2. Install **debootstrap** ... 89
D.3.3. Run **debootstrap** ... 90
D.3.4. Configure The Base System ... 90
D.3.4.1. Mount Partitions ... 90
D.3.4.2. Setting Timezone ... 91
D.3.4.3. Configure Networking ... 91
D.3.4.4. Configure Apt ... 93
D.3.4.5. Configure Locales and Keyboard ... 93
D.3.5. Install a Kernel ... 93
D.3.6. Set up the Boot Loader ... 94
D.3.7. Finishing touches ... 94
D.4. Installing Debian GNU/Linux using PPP over Ethernet (PPPoE) ... 94
E. Administrivia ... 96
E.1. About This Document ... 96
E.2. Contributing to This Document ... 96
E.3. Major Contributions ... 96
E.4. Trademark Acknowledgement ... 97
F. GNU General Public License ... 98
F.1. Preamble ... 98
F.2. GNU GENERAL PUBLIC LICENSE ... 98
F.3. How to Apply These Terms to Your New Programs ... 102

List of Tables

3-1. Hardware Information Needed for an Install 13
3-2. Recommended Minimum System Requirements 15

Installing Debian GNU/Linux 4.0 For hppa

We are delighted that you have decided to try Debian, and are sure that you will find that Debian's GNU/Linux distribution is unique. Debian GNU/Linux brings together high-quality free software from around the world, integrating it into a coherent whole. We believe that you will find that the result is truly more than the sum of the parts.

We understand that many of you want to install Debian without reading this manual, and the Debian installer is designed to make this possible. If you don't have time to read the whole Installation Guide right now, we recommend that you read the Installation Howto, which will walk you through the basic installation process, and links to the manual for more advanced topics or for when things go wrong. The Installation Howto can be found in Appendix A.

With that said, we hope that you have the time to read most of this manual, and doing so will lead to a more informed and likely more successful installation experience.

Chapter 1. Welcome to Debian

This chapter provides an overview of the Debian Project and Debian GNU/Linux. If you already know about the Debian Project's history and the Debian GNU/Linux distribution, feel free to skip to the next chapter.

1.1. What is Debian?

Debian is an all-volunteer organization dedicated to developing free software and promoting the ideals of the Free Software community. The Debian Project began in 1993, when Ian Murdock issued an open invitation to software developers to contribute to a complete and coherent software distribution based on the relatively new Linux kernel. That relatively small band of dedicated enthusiasts, originally funded by the Free Software Foundation (http://www.fsf.org/) and influenced by the GNU (http://www.gnu.org/gnu/the-gnu-project.html) philosophy, has grown over the years into an organization of around 1010 *Debian Developers*.

Debian Developers are involved in a variety of activities, including Web (http://www.debian.org/) and FTP (ftp://ftp.debian.org/) site administration, graphic design, legal analysis of software licenses, writing documentation, and, of course, maintaining software packages.

In the interest of communicating our philosophy and attracting developers who believe in the principles that Debian stands for, the Debian Project has published a number of documents that outline our values and serve as guides to what it means to be a Debian Developer:

- The Debian Social Contract (http://www.debian.org/social_contract) is a statement of Debian's commitments to the Free Software Community. Anyone who agrees to abide to the Social Contract may become a maintainer (http://www.debian.org/doc/maint-guide/). Any maintainer can introduce new software into Debian — provided that the software meets our criteria for being free, and the package follows our quality standards.
- The Debian Free Software Guidelines (http://www.debian.org/social_contract#guidelines) are a clear and concise statement of Debian's criteria for free software. The DFSG is a very influential document in the Free Software Movement, and was the foundation of the The Open Source Definition (http://opensource.org/docs/definition_plain.html).
- The Debian Policy Manual (http://www.debian.org/doc/debian-policy/) is an extensive specification of the Debian Project's standards of quality.

Debian developers are also involved in a number of other projects; some specific to Debian, others involving some or all of the Linux community. Some examples include:

- The Linux Standard Base (http://www.linuxbase.org/) (LSB) is a project aimed at standardizing the basic GNU/Linux system, which will enable third-party software and hardware developers to easily design programs and device drivers for Linux-in-general, rather than for a specific GNU/Linux distribution.
- The Filesystem Hierarchy Standard (http://www.pathname.com/fhs/) (FHS) is an effort to standardize the layout of the Linux file system. The FHS will allow software developers to concentrate their efforts on designing programs, without having to worry about how the package will be installed in different GNU/Linux distributions.

- Debian Jr. (http://www.debian.org/devel/debian-jr/) is an internal project, aimed at making sure Debian has something to offer to our youngest users.

For more general information about Debian, see the Debian FAQ (http://www.debian.org/doc/FAQ/).

1.2. What is GNU/Linux?

Linux is an operating system: a series of programs that let you interact with your computer and run other programs.

An operating system consists of various fundamental programs which are needed by your computer so that it can communicate and receive instructions from users; read and write data to hard disks, tapes, and printers; control the use of memory; and run other software. The most important part of an operating system is the kernel. In a GNU/Linux system, Linux is the kernel component. The rest of the system consists of other programs, many of which were written by or for the GNU Project. Because the Linux kernel alone does not form a working operating system, we prefer to use the term "GNU/Linux" to refer to systems that many people casually refer to as "Linux".

Linux is modelled on the Unix operating system. From the start, Linux was designed to be a multi-tasking, multi-user system. These facts are enough to make Linux different from other well-known operating systems. However, Linux is even more different than you might imagine. In contrast to other operating systems, nobody owns Linux. Much of its development is done by unpaid volunteers.

Development of what later became GNU/Linux began in 1984, when the Free Software Foundation (http://www.fsf.org/) began development of a free Unix-like operating system called GNU.

The GNU Project (http://www.gnu.org/) has developed a comprehensive set of free software tools for use with Unix™ and Unix-like operating systems such as Linux. These tools enable users to perform tasks ranging from the mundane (such as copying or removing files from the system) to the arcane (such as writing and compiling programs or doing sophisticated editing in a variety of document formats).

While many groups and individuals have contributed to Linux, the largest single contributor is still the Free Software Foundation, which created not only most of the tools used in Linux, but also the philosophy and the community that made Linux possible.

The Linux kernel (http://www.kernel.org/) first appeared in 1991, when a Finnish computing science student named Linus Torvalds announced an early version of a replacement kernel for Minix to the Usenet newsgroup **`comp.os.minix`**. See Linux International's Linux History Page (http://www.li.org/linuxhistory.php).

Linus Torvalds continues to coordinate the work of several hundred developers with the help of a few trusty deputies. An excellent weekly summary of discussions on the **`linux-kernel`** mailing list is Kernel Traffic (http://www.kerneltraffic.org/kernel-traffic/index.html). More information about the **`linux-kernel`** mailing list can be found on the linux-kernel mailing list FAQ (http://www.tux.org/lkml/).

Linux users have immense freedom of choice in their software. For example, Linux users can choose from a dozen different command line shells and several graphical desktops. This selection is often bewildering to users of other operating systems, who are not used to thinking of the command line or desktop as something that they can change.

Linux is also less likely to crash, better able to run more than one program at the same time, and more secure than many operating systems. With these advantages, Linux is the fastest growing operating

system in the server market. More recently, Linux has begun to be popular among home and business users as well.

1.3. What is Debian GNU/Linux?

The combination of Debian's philosophy and methodology and the GNU tools, the Linux kernel, and other important free software, form a unique software distribution called Debian GNU/Linux. This distribution is made up of a large number of software *packages*. Each package in the distribution contains executables, scripts, documentation, and configuration information, and has a *maintainer* who is primarily responsible for keeping the package up-to-date, tracking bug reports, and communicating with the upstream author(s) of the packaged software. Our extremely large user base, combined with our bug tracking system ensures that problems are found and fixed quickly.

Debian's attention to detail allows us to produce a high-quality, stable, and scalable distribution. Installations can be easily configured to serve many roles, from stripped-down firewalls to desktop scientific workstations to high-end network servers.

Debian is especially popular among advanced users because of its technical excellence and its deep commitment to the needs and expectations of the Linux community. Debian also introduced many features to Linux that are now commonplace.

For example, Debian was the first Linux distribution to include a package management system for easy installation and removal of software. It was also the first Linux distribution that could be upgraded without requiring reinstallation.

Debian continues to be a leader in Linux development. Its development process is an example of just how well the Open Source development model can work — even for very complex tasks such as building and maintaining a complete operating system.

The feature that most distinguishes Debian from other Linux distributions is its package management system. These tools give the administrator of a Debian system complete control over the packages installed on that system, including the ability to install a single package or automatically update the entire operating system. Individual packages can also be protected from being updated. You can even tell the package management system about software you have compiled yourself and what dependencies it fulfills.

To protect your system against "Trojan horses" and other malevolent software, Debian's servers verify that uploaded packages come from their registered Debian maintainers. Debian packagers also take great care to configure their packages in a secure manner. When security problems in shipped packages do appear, fixes are usually available very quickly. With Debian's simple update options, security fixes can be downloaded and installed automatically across the Internet.

The primary, and best, method of getting support for your Debian GNU/Linux system and communicating with Debian Developers is through the many mailing lists maintained by the Debian Project (there are more than 215 at this writing). The easiest way to subscribe to one or more of these lists is visit Debian's mailing list subscription page (http://www.debian.org/MailingLists/subscribe) and fill out the form you'll find there.

1.4. Getting Debian

For information on how to download Debian GNU/Linux from the Internet or from whom official Debian CDs can be purchased, see the distribution web page (http://www.debian.org/distrib/). The list

of Debian mirrors (http://www.debian.org/distrib/ftplist) contains a full set of official Debian mirrors, so you can easily find the nearest one.

Debian can be upgraded after installation very easily. The installation procedure will help set up the system so that you can make those upgrades once installation is complete, if need be.

1.5. Getting the Newest Version of This Document

This document is constantly being revised. Be sure to check the Debian 4.0 pages (http://www.debian.org/releases/etch/) for any last-minute information about the 4.0 release of the Debian GNU/Linux system. Updated versions of this installation manual are also available from the official Install Manual pages (http://www.debian.org/releases/etch/hppa/).

1.6. Organization of This Document

This document is meant to serve as a manual for first-time Debian users. It tries to make as few assumptions as possible about your level of expertise. However, we do assume that you have a general understanding of how the hardware in your computer works.

Expert users may also find interesting reference information in this document, including minimum installation sizes, details about the hardware supported by the Debian installation system, and so on. We encourage expert users to jump around in the document.

In general, this manual is arranged in a linear fashion, walking you through the installation process from start to finish. Here are the steps in installing Debian GNU/Linux, and the sections of this document which correlate with each step:

1. Determine whether your hardware meets the requirements for using the installation system, in Chapter 2.
2. Backup your system, perform any necessary planning and hardware configuration prior to installing Debian, in Chapter 3. If you are preparing a multi-boot system, you may need to create partition-able space on your hard disk for Debian to use.
3. In Chapter 4, you will obtain the necessary installation files for your method of installation.
4. Chapter 5 describes booting into the installation system. This chapter also discusses troubleshooting procedures in case you have problems with this step.
5. Perform the actual installation according to Chapter 6. This involves choosing your language, configuring peripheral driver modules, configuring your network connection, so that remaining installation files can be obtained directly from a Debian server (if you are not installing from a CD), partitioning your hard drives and installation of a base system, then selection and installation of tasks. (Some background about setting up the partitions for your Debian system is explained in Appendix C.)
6. Boot into your newly installed base system, from Chapter 7.

Once you've got your system installed, you can read Chapter 8. That chapter explains where to look to find more information about Unix and Debian, and how to replace your kernel.

Finally, information about this document and how to contribute to it may be found in Appendix E.

1.7. About Copyrights and Software Licenses

We're sure that you've read some of the licenses that come with most commercial software — they usually say that you can only use one copy of the software on a single computer. This system's license isn't like that at all. We encourage you to put a copy of on every computer in your school or place of business. Lend your installation media to your friends and help them install it on their computers! You can even make thousands of copies and *sell* them — albeit with a few restrictions. Your freedom to install and use the system comes directly from Debian being based on *free software*.

Calling software *free* doesn't mean that the software isn't copyrighted, and it doesn't mean that CDs containing that software must be distributed at no charge. Free software, in part, means that the licenses of individual programs do not require you to pay for the privilege of distributing or using those programs. Free software also means that not only may anyone extend, adapt, and modify the software, but that they may distribute the results of their work as well.

> **Note:** The Debian project, as a pragmatic concession to its users, does make some packages available that do not meet our criteria for being free. These packages are not part of the official distribution, however, and are only available from the **`contrib`** or **`non-free`** areas of Debian mirrors or on third-party CD-ROMs; see the Debian FAQ (http://www.debian.org/doc/FAQ/), under "The Debian FTP archives", for more information about the layout and contents of the archives.

Many of the programs in the system are licensed under the *GNU General Public License*, often simply referred to as "the GPL". The GPL requires you to make the *source code* of the programs available whenever you distribute a binary copy of the program; that provision of the license ensures that any user will be able to modify the software. Because of this provision, the source code[1] for all such programs is available in the Debian system.

There are several other forms of copyright statements and software licenses used on the programs in Debian. You can find the copyrights and licenses for every package installed on your system by looking in the file `/usr/share/doc/`*`package-name`*`/copyright` once you've installed a package on your system.

For more information about licenses and how Debian determines whether software is free enough to be included in the main distribution, see the Debian Free Software Guidelines (http://www.debian.org/social_contract#guidelines).

The most important legal notice is that this software comes with *no warranties*. The programmers who have created this software have done so for the benefit of the community. No guarantee is made as to the suitability of the software for any given purpose. However, since the software is free, you are empowered to modify that software to suit your needs — and to enjoy the benefits of the changes made by others who have extended the software in this way.

1. For information on how to locate, unpack, and build binaries from Debian source packages, see the Debian FAQ (http://www.debian.org/doc/FAQ/), under "Basics of the Debian Package Management System".

Chapter 2. System Requirements

This section contains information about what hardware you need to get started with Debian. You will also find links to further information about hardware supported by GNU and Linux.

2.1. Supported Hardware

Debian does not impose hardware requirements beyond the requirements of the Linux kernel and the GNU tool-sets. Therefore, any architecture or platform to which the Linux kernel, libc, **gcc**, etc. have been ported, and for which a Debian port exists, can run Debian. Please refer to the Ports pages at http://www.debian.org/ports/hppa/ for more details on PA-RISC architecture systems which have been tested with Debian.

Rather than attempting to describe all the different hardware configurations which are supported for PA-RISC, this section contains general information and pointers to where additional information can be found.

2.1.1. Supported Architectures

Debian 4.0 supports twelve major architectures and several variations of each architecture known as “flavors”.

Architecture	Debian Designation	Subarchitecture	Flavor
Intel x86-based	i386		
AMD64 & Intel EM64T	amd64		
DEC Alpha	alpha		
ARM and StrongARM	arm	Netwinder and CATS	netwinder
		Intel IOP32x	iop32x
		Intel IXP4xx	ixp4xx
		RiscPC	rpc
HP PA-RISC	hppa	PA-RISC 1.1	32
		PA-RISC 2.0	64
Intel IA-64	ia64		
MIPS (big endian)	mips	SGI IP22 (Indy/Indigo 2)	r4k-ip22
		SGI IP32 (O2)	r5k-ip32
		Broadcom BCM91250A (SWARM)	sb1-bcm91250a
		Broadcom BCM91480B (BigSur)	sb1a-bcm91480b
MIPS (little endian)	mipsel	Cobalt	cobalt

Architecture	Debian Designation	Subarchitecture	Flavor
		DECstation	r4k-kn04
			r3k-kn02
		Broadcom BCM91250A (SWARM)	sb1-bcm91250a
		Broadcom BCM91480B (BigSur)	sb1a-bcm91480b
Motorola 680x0	m68k	Atari	atari
		Amiga	amiga
		68k Macintosh	mac
		VME	bvme6000
			mvme147
			mvme16x
IBM/Motorola PowerPC	powerpc	CHRP	chrp
		PowerMac	pmac
		PReP	prep
Sun SPARC	sparc	sun4m	sparc32
		sun4u	sparc64
		sun4v	
IBM S/390	s390	IPL from VM-reader and DASD	generic
		IPL from tape	tape

This document covers installation for the *PA-RISC* architecture. If you are looking for information on any of the other Debian-supported architectures take a look at the Debian-Ports (http://www.debian.org/ports/) pages.

2.1.2. CPU, Main Boards, and Video Support

There are two major support *hppa* flavors: PA-RISC 1.1 and PA-RISC 2.0. The PA-RISC 1.1 architecture is targeted at 32-bit processors whereas the 2.0 architecture is targeted to the 64-bit processors. Some systems are able to run either kernel. In both cases, the userland is 32-bit. There is the possibility of a 64-bit userland in the future.

2.1.3. Graphics Card Support

Debian's support for graphical interfaces is determined by the underlying support found in X.Org's X11 system. Most AGP, PCI and PCIe video cards work under X.Org. Details on supported graphics

buses, cards, monitors, and pointing devices can be found at http://xorg.freedesktop.org/. Debian 4.0 ships with X.Org version 7.1.

2.1.4. Multiple Processors

Multiprocessor support — also called "symmetric multiprocessing" or SMP — is available for this architecture. However, the standard Debian 4.0 kernel image does not support SMP. This should not prevent installation, since the standard, non-SMP kernel should boot on SMP systems; the kernel will simply use the first CPU.

In order to take advantage of multiple processors, you'll have to replace the standard Debian kernel. You can find a discussion of how to do this in Section 8.6. At this time (kernel version 2.6.18) the way you enable SMP is to select "Symmetric multi-processing support" in the "Processor type and features" section of the kernel config.

2.2. Installation Media

This section will help you determine which different media types you can use to install Debian. For example, if you have a floppy disk drive on your machine, it can be used to install Debian. There is a whole chapter devoted to media, Chapter 4, which lists the advantages and disadvantages of each media type. You may want to refer back to this page once you reach that section.

2.2.1. CD-ROM/DVD-ROM

> **Note:** Whenever you see "CD-ROM" in this manual, it applies to both CD-ROMs and DVD-ROMs, because both technologies are really the same from the operating system's point of view, except for some very old nonstandard CD-ROM drives which are neither SCSI nor IDE/ATAPI.

CD-ROM based installation is supported for some architectures. On machines which support bootable CD-ROMs, you should be able to do a completely floppy-less installation. Even if your system doesn't support booting from a CD-ROM, you can use the CD-ROM in conjunction with the other techniques to install your system, once you've booted up by other means; see Chapter 5.

2.2.2. Hard Disk

Booting the installation system directly from a hard disk is another option for many architectures. This will require some other operating system to load the installer onto the hard disk.

2.2.3. Network

The network can be used during the installation to retrieve files needed for the installation. Whether the network is used or not depends on the installation method you choose and your answers to certain questions that will be asked during the installation. The installation system supports most types of network connections (including PPPoE, but not ISDN or PPP), via either HTTP or FTP. After the installation is completed, you can also configure your system to use ISDN and PPP.

You can also *boot* the installation system over the network.

Diskless installation, using network booting from a local area network and NFS-mounting of all local filesystems, is another option.

2.2.4. Un*x or GNU system

If you are running another Unix-like system, you could use it to install Debian GNU/Linux without using the `debian-installer` described in the rest of this manual. This kind of install may be useful for users with otherwise unsupported hardware or on hosts which can't afford downtime. If you are interested in this technique, skip to the Section D.3.

2.2.5. Supported Storage Systems

The Debian boot disks contain a kernel which is built to maximize the number of systems it runs on. Unfortunately, this makes for a larger kernel, which includes many drivers that won't be used for your machine (see Section 8.6 to learn how to build your own kernel). Support for the widest possible range of devices is desirable in general, to ensure that Debian can be installed on the widest array of hardware.

Any storage system supported by the Linux kernel is also supported by the boot system. Note that the current Linux kernel does not support the floppy drive.

2.3. Peripherals and Other Hardware

Linux supports a large variety of hardware devices such as mice, printers, scanners, PCMCIA and USB devices. However, most of these devices are not required while installing the system.

2.4. Purchasing Hardware Specifically for GNU/Linux

There are several vendors, who ship systems with Debian or other distributions of GNU/Linux pre-installed (http://www.debian.org/distrib/pre-installed). You might pay more for the privilege, but it does buy a level of peace of mind, since you can be sure that the hardware is well-supported by GNU/Linux.

Whether or not you are purchasing a system with Linux bundled, or even a used system, it is still important to check that your hardware is supported by the Linux kernel. Check if your hardware is listed in the references found above. Let your salesperson (if any) know that you're shopping for a Linux system. Support Linux-friendly hardware vendors.

2.4.1. Avoid Proprietary or Closed Hardware

Some hardware manufacturers simply won't tell us how to write drivers for their hardware. Others won't allow us access to the documentation without a non-disclosure agreement that would prevent us from releasing the Linux source code.

Since we haven't been granted access to the documentation on these devices, they simply won't work under Linux. You can help by asking the manufacturers of such hardware to release the documentation. If enough people ask, they will realize that the free software community is an important market.

2.5. Memory and Disk Space Requirements

You must have at least 32MB of memory and 500MB of hard disk space. Note that these are really bare-minimum numbers. For more realistic figures, see Section 3.4.

2.6. Network Connectivity Hardware

Almost any network interface card (NIC) supported by the Linux kernel should also be supported by the installation system; modular drivers should normally be loaded automatically.

2.6.1. Drivers Requiring Firmware

The installation system currently does not support retrieving firmware. This means that any network cards that use a driver that requires firmware to be loaded, is not supported by default.

If there is no other NIC you can use during the installation, it is still possible to install Debian GNU/Linux using a full CD-ROM or DVD image. Select the option to not configure a network and install using only the packages available from the CD/DVD. You can then install the driver and firmware you need after the installation is completed (after the reboot) and configure your network manually. Note that the firmware may be packaged separately from the driver and may not be available in the "main" section of the Debian GNU/Linux archive.

If the driver itself *is* supported, you may also be able to use the NIC during installation by copying the firmware from some medium to `/usr/lib/hotplug/firmware`. Don't forget to also copy the firmware to that location for the installed system before the reboot at the end of the installation.

Chapter 3. Before Installing Debian GNU/Linux

This chapter deals with the preparation for installing Debian before you even boot the installer. This includes backing up your data, gathering information about your hardware, and locating any necessary information.

3.1. Overview of the Installation Process

First, just a note about re-installations. With Debian, a circumstance that will require a complete re-installation of your system is very rare; perhaps mechanical failure of the hard disk would be the most common case.

Many common operating systems may require a complete installation to be performed when critical failures take place or for upgrades to new OS versions. Even if a completely new installation isn't required, often the programs you use must be re-installed to operate properly in the new OS.

Under Debian GNU/Linux, it is much more likely that your OS can be repaired rather than replaced if things go wrong. Upgrades never require a wholesale installation; you can always upgrade in-place. And the programs are almost always compatible with successive OS releases. If a new program version requires newer supporting software, the Debian packaging system ensures that all the necessary software is automatically identified and installed. The point is, much effort has been put into avoiding the need for re-installation, so think of it as your very last option. The installer is *not* designed to re-install over an existing system.

Here's a road map for the steps you will take during the installation process.

1. Back up any existing data or documents on the hard disk where you plan to install.
2. Gather information about your computer and any needed documentation, before starting the installation.
3. Create partition-able space for Debian on your hard disk.
4. Locate and/or download the installer software and any specialized driver files your machine requires (except Debian CD users).
5. Set up boot tapes/floppies/USB sticks, or place boot files (most Debian CD users can boot from one of the CDs).
6. Boot the installation system.
7. Select the installation language.
8. Activate the ethernet network connection, if available.
9. Create and mount the partitions on which Debian will be installed.
10. Watch the automatic download/install/setup of the *base system*.
11. Install a *boot loader* which can start up Debian GNU/Linux and/or your existing system.
12. Load the newly installed system for the first time.

If you have problems during the installation, it helps to know which packages are involved in which steps. Introducing the leading software actors in this installation drama:

The installer software, `debian-installer`, is the primary concern of this manual. It detects hardware and loads appropriate drivers, uses `dhcp-client` to set up the network connection, runs `debootstrap` to install the base system packages, and runs `tasksel` to allow you to install certain additional software. Many more actors play smaller parts in this process, but `debian-installer` has completed its task when you load the new system for the first time.

To tune the system to your needs, `tasksel` allows you to choose to install various predefined bundles of software like a Web server or a Desktop environment.

One important option during the installation is whether or not to install a graphical desktop environment, consisting of the X Window System and one of the available graphical desktop environments. If you choose not to select the "Desktop environment" task, you will only have a relatively basic, command line driven system. Installing the Desktop environment task is optional because it requires a fairly large amount of disk space, and because many Debian GNU/Linux systems are servers which don't really have any need for a graphical user interface to do their job.

Just be aware that the X Window System is completely separate from `debian-installer`, and in fact is much more complicated. Installation and troubleshooting of the X Window System is not within the scope of this manual.

3.2. Back Up Your Existing Data!

Before you start, make sure to back up every file that is now on your system. If this is the first time a non-native operating system has been installed on your computer, it's quite likely you will need to re-partition your disk to make room for Debian GNU/Linux. Anytime you partition your disk, you run a risk of losing everything on the disk, no matter what program you use to do it. The programs used in installation are quite reliable and most have seen years of use; but they are also quite powerful and a false move can cost you. Even after backing up, be careful and think about your answers and actions. Two minutes of thinking can save hours of unnecessary work.

If you are creating a multi-boot system, make sure that you have the distribution media of any other present operating systems on hand. Especially if you repartition your boot drive, you might find that you have to reinstall your operating system's boot loader, or in many cases the whole operating system itself and all files on the affected partitions.

3.3. Information You Will Need

3.3.1. Documentation

3.3.1.1. Installation Manual

The document you are now reading, which is the official version of the Installation Guide for the etch release of Debian; available in various formats and translations (http://www.debian.org/releases/etch//installmanual).

3.3.1.2. Hardware documentation

Often contains useful information on configuring or using your hardware.

3.3.2. Finding Sources of Hardware Information

In many cases, the installer will be able to automatically detect your hardware. But to be prepared, we do recommend familiarizing yourself with your hardware before the install.

Hardware information can be gathered from:

- The manuals that come with each piece of hardware.
- The BIOS setup screens of your computer. You can view these screens when you start your computer by pressing a combination of keys. Check your manual for the combination. Often, it is the **Delete** key.
- The cases and boxes for each piece of hardware.
- System commands or tools in another operating system, including file manager displays. This source is especially useful for information about RAM and hard drive memory.
- Your system administrator or Internet Service Provider. These sources can tell you the settings you need to set up your networking and e-mail.

Table 3-1. Hardware Information Needed for an Install

Hardware	Information You Might Need
Hard Drives	How many you have.
	Their order on the system.
	Whether IDE or SCSI (most computers are IDE).
	Available free space.
	Partitions.
	Partitions where other operating systems are installed.
Monitor	Model and manufacturer.
	Resolutions supported.
	Horizontal refresh rate.
	Vertical refresh rate.
	Color depth (number of colors) supported.
	Screen size.
Mouse	Type: serial, PS/2, or USB.
	Port.
	Manufacturer.
	Number of buttons.
Network	Model and manufacturer.
	Type of adapter.
Printer	Model and manufacturer.
	Printing resolutions supported.

Hardware	Information You Might Need
Video Card	Model and manufacturer.
	Video RAM available.
	Resolutions and color depths supported (these should be checked against your monitor's capabilities).

3.3.3. Hardware Compatibility

Many brand name products work without trouble on Linux. Moreover, hardware support in Linux is improving daily. However, Linux still does not run as many different types of hardware as some operating systems.

You can check hardware compatibility by:

- Checking manufacturers' web sites for new drivers.
- Looking at web sites or manuals for information about emulation. Lesser known brands can sometimes use the drivers or settings for better-known ones.
- Checking hardware compatibility lists for Linux on web sites dedicated to your architecture.
- Searching the Internet for other users' experiences.

3.3.4. Network Settings

If your computer is connected to a network 24 hours a day (i.e., an Ethernet or equivalent connection — not a PPP connection), you should ask your network's system administrator for this information.

- Your host name (you may be able to decide this on your own).
- Your domain name.
- Your computer's IP address.
- The netmask to use with your network.
- The IP address of the default gateway system you should route to, if your network *has* a gateway.
- The system on your network that you should use as a DNS (Domain Name Service) server.

On the other hand, if your administrator tells you that a DHCP server is available and is recommended, then you don't need this information because the DHCP server will provide it directly to your computer during the installation process.

If you use a wireless network, you should also find out:

- ESSID of your wireless network.
- WEP security key (if applicable).

3.4. Meeting Minimum Hardware Requirements

Once you have gathered information about your computer's hardware, check that your hardware will let you do the type of installation that you want to do.

Depending on your needs, you might manage with less than some of the recommended hardware listed in the table below. However, most users risk being frustrated if they ignore these suggestions.

Table 3-2. Recommended Minimum System Requirements

Install Type	RAM (minimal)	RAM (recommended)	Hard Drive
No desktop	64 megabytes	256 megabytes	1 gigabyte
With Desktop	64 megabytes	512 megabytes	5 gigabyte

The actual minimum memory requirements are a lot less then the numbers listed in this table. Depending on the architecture, it is possible to install Debian with as little as 20MB (for s390) to 48MB (for i386 and amd64). The same goes for the disk space requirements, especially if you pick and choose which applications to install; see Section D.2 for additional information on disk space requirements.

It is possible to run a graphical desktop environment on older or low-end systems, but in that case it is recommended to install a window manager that is less resource-hungry than those of the GNOME or KDE desktop environments; alternatives include `xfce4`, `icewm` and `wmaker`, but there are others to choose from.

It is practically impossible to give general memory or disk space requirements for server installations as those very much depend on what the server is to be used for.

Remember that these sizes don't include all the other materials which are usually to be found, such as user files, mail, and data. It is always best to be generous when considering the space for your own files and data.

Disk space required for the smooth operation of the Debian GNU/Linux system itself is taken into account in these recommended system requirements. Notably, the `/var` partition contains a lot of state information specific to Debian in addition to its regular contents, like logfiles. The **dpkg** files (with information on all installed packages) can easily consume 40MB. Also, **apt-get** puts downloaded packages here before they are installed. You should usually allocate at least 200MB for `/var`, and a lot more if you install a graphical desktop environment.

3.5. Pre-Partitioning for Multi-Boot Systems

Partitioning your disk simply refers to the act of breaking up your disk into sections. Each section is then independent of the others. It's roughly equivalent to putting up walls inside a house; if you add furniture to one room it doesn't affect any other room.

If you already have an operating system on your system and want to stick Linux on the same disk, you will need to repartition the disk. Debian requires its own hard disk partitions. It cannot be installed on Windows or MacOS partitions. It may be able to share some partitions with other Linux systems, but that's not covered here. At the very least you will need a dedicated partition for the Debian root.

You can find information about your current partition setup by using a partitioning tool for your current operating system. Partitioning tools always provide a way to show existing partitions without making changes.

In general, changing a partition with a file system already on it will destroy any information there. Thus you should always make backups before doing any repartitioning. Using the analogy of the house, you would probably want to move all the furniture out of the way before moving a wall or you risk destroying it.

If your computer has more than one hard disk, you may want to dedicate one of the hard disks completely to Debian. If so, you don't need to partition that disk before booting the installation system; the installer's included partitioning program can handle the job nicely.

If your machine has only one hard disk, and you would like to completely replace the current operating system with Debian GNU/Linux, you also can wait to partition as part of the installation process (Section 6.3.2.1), after you have booted the installation system. However this only works if you plan to boot the installer system from tapes, CD-ROM or files on a connected machine. Consider: if you boot from files placed on the hard disk, and then partition that same hard disk within the installation system, thus erasing the boot files, you'd better hope the installation is successful the first time around. At the least in this case, you should have some alternate means of reviving your machine like the original system's installation tapes or CDs.

If your machine already has multiple partitions, and enough space can be provided by deleting and replacing one or more of them, then you too can wait and use the Debian installer's partitioning program. You should still read through the material below, because there may be special circumstances like the order of the existing partitions within the partition map, that force you to partition before installing anyway.

If none of the above apply, you'll need to partition your hard disk before starting the installation to create partition-able space for Debian. If some of the partitions will be owned by other operating systems, you should create those partitions using native operating system partitioning programs. We recommend that you do *not* attempt to create partitions for Debian GNU/Linux using another operating system's tools. Instead, you should just create the native operating system's partitions you will want to retain.

If you are going to install more than one operating system on the same machine, you should install all other system(s) before proceeding with Linux installation. Windows and other OS installations may destroy your ability to start Linux, or encourage you to reformat non-native partitions.

You can recover from these actions or avoid them, but installing the native system first saves you trouble.

If you currently have one hard disk with one partition (a common setup for desktop computers), and you want to multi-boot the native operating system and Debian, you will need to:

1. Back up everything on the computer.
2. Boot from the native operating system installer media such as CD-ROM or tapes.
3. Use the native partitioning tools to create native system partition(s). Leave either a place holder partition or free space for Debian GNU/Linux.
4. Install the native operating system on its new partition.
5. Boot back into the native system to verify everything's OK, and to download the Debian installer boot files.
6. Boot the Debian installer to continue installing Debian.

3.6. Pre-Installation Hardware and Operating System Setup

This section will walk you through pre-installation hardware setup, if any, that you will need to do prior to installing Debian. Generally, this involves checking and possibly changing firmware settings for your system. The “firmware” is the core software used by the hardware; it is most critically invoked during the bootstrap process (after power-up). Known hardware issues affecting the reliability of Debian GNU/Linux on your system are also highlighted.

Chapter 4. Obtaining System Installation Media

4.1. Official Debian GNU/Linux CD-ROM Sets

By far the easiest way to install Debian GNU/Linux is from an Official Debian CD-ROM Set. You can buy a set from a vendor (see the CD vendors page (http://www.debian.org/CD/vendors/)). You may also download the CD-ROM images from a Debian mirror and make your own set, if you have a fast network connection and a CD burner (see the Debian CD page (http://www.debian.org/CD/) for detailed instructions). If you have a Debian CD set and CDs are bootable on your machine, you can skip right to Chapter 5; much effort has been expended to ensure the files most people need are there on the CD. Although a full set of binary packages requires several CDs, it is unlikely you will need packages on the third CD and above. You may also consider using the DVD version, which saves a lot of space on your shelf and you avoid the CD shuffling marathon.

If your machine doesn't support CD booting, but you do have a CD set, you can use an alternative strategy such as net boot, or manually loading the kernel from the CD to initially boot the system installer. The files you need for booting by another means are also on the CD; the Debian network archive and CD folder organization are identical. So when archive file paths are given below for particular files you need for booting, look for those files in the same directories and subdirectories on your CD.

Once the installer is booted, it will be able to obtain all the other files it needs from the CD.

If you don't have a CD set, then you will need to download the installer system files and place them on the a connected computer so they can be used to boot the installer.

4.2. Downloading Files from Debian Mirrors

To find the nearest (and thus probably the fastest) mirror, see the list of Debian mirrors (http://www.debian.org/distrib/ftplist).

When downloading files from a Debian mirror, be sure to download the files in *binary* mode, not text or automatic mode.

4.2.1. Where to Find Installation Images

The installation images are located on each Debian mirror in the directory debian/dists/etch/main/installer-hppa/current/images/ (http://http.us.debian.org/debian/dists/etch/main/installer-hppa/current//images) — the MANIFEST (http://http.us.debian.org/debian/dists/etch/main/installer-hppa/current//images/MANIFEST) lists each image and its purpose.

4.3. Preparing Files for TFTP Net Booting

If your machine is connected to a local area network, you may be able to boot it over the network from

another machine, using TFTP. If you intend to boot the installation system from another machine, the boot files will need to be placed in specific locations on that machine, and the machine configured to support booting of your specific machine.

You need to setup a TFTP server, and for many machines a DHCP server, or BOOTP server.

BOOTP is an IP protocol that informs a computer of its IP address and where on the network to obtain a boot image. The DHCP (Dynamic Host Configuration Protocol) is a more flexible, backwards-compatible extension of BOOTP. Some systems can only be configured via DHCP.

Some older HPPA machines (e.g. 715/75) use RBOOTD rather than BOOTP. There is an `rbootd` package available in Debian.

The Trivial File Transfer Protocol (TFTP) is used to serve the boot image to the client. Theoretically, any server, on any platform, which implements these protocols, may be used. In the examples in this section, we shall provide commands for SunOS 4.x, SunOS 5.x (a.k.a. Solaris), and GNU/Linux.

4.3.1. Setting up a BOOTP server

There are two BOOTP servers available for GNU/Linux. The first is CMU **bootpd**. The other is actually a DHCP server: ISC **dhcpd**. In Debian GNU/Linux these are contained in the `bootp` and `dhcp3-server` packages respectively.

To use CMU **bootpd**, you must first uncomment (or add) the relevant line in `/etc/inetd.conf`. On Debian GNU/Linux, you can run **`update-inetd --enable bootps`**, then **`/etc/init.d/inetd reload`** to do so. Just in case your BOOTP server does not run Debian, the line in question should look like:

```
bootps  dgram  udp  wait  root  /usr/sbin/bootpd  bootpd -i -t 120
```

Now, you must create an `/etc/bootptab` file. This has the same sort of familiar and cryptic format as the good old BSD `printcap`, `termcap`, and `disktab` files. See the `bootptab` manual page for more information. For CMU **bootpd**, you will need to know the hardware (MAC) address of the client. Here is an example `/etc/bootptab`:

```
client:\
  hd=/tftpboot:\
  bf=tftpboot.img:\
  ip=192.168.1.90:\
  sm=255.255.255.0:\
  sa=192.168.1.1:\
  ha=0123456789AB:
```

You will need to change at least the “ha” option, which specifies the hardware address of the client. The “bf” option specifies the file a client should retrieve via TFTP; see Section 4.3.4 for more details.

By contrast, setting up BOOTP with ISC **dhcpd** is really easy, because it treats BOOTP clients as a moderately special case of DHCP clients. Some architectures require a complex configuration for booting clients via BOOTP. If yours is one of those, read the section Section 4.3.2. In that case, you will probably be able to get away with simply adding the **`allow bootp`** directive to the configuration block for the subnet containing the client, and restart **dhcpd** with **`/etc/init.d/dhcpd3-server restart`**.

4.3.2. Setting up a DHCP server

One free software DHCP server is ISC **dhcpd**. For Debian GNU/Linux, the `dhcp3-server` package is recommended. Here is a sample configuration file for it (see `/etc/dhcp3/dhcpd.conf`):

```
option domain-name "example.com";
option domain-name-servers ns1.example.com;
option subnet-mask 255.255.255.0;
default-lease-time 600;
max-lease-time 7200;
server-name "servername";

subnet 192.168.1.0 netmask 255.255.255.0 {
  range 192.168.1.200 192.168.1.253;
  option routers 192.168.1.1;
}

host clientname {
  filename "/tftpboot/tftpboot.img";
  server-name "servername";
  next-server servername;
  hardware ethernet 01:23:45:67:89:AB;
  fixed-address 192.168.1.90;
}
```

In this example, there is one server *`servername`* which performs all of the work of DHCP server, TFTP server, and network gateway. You will almost certainly need to change the domain-name options, as well as the server name and client hardware address. The *`filename`* option should be the name of the file which will be retrieved via TFTP.

After you have edited the **dhcpd** configuration file, restart it with **`/etc/init.d/dhcpd3-server restart`**.

4.3.3. Enabling the TFTP Server

To get the TFTP server ready to go, you should first make sure that **tftpd** is enabled. This is usually enabled by having something like the following line in `/etc/inetd.conf`:

```
tftp dgram udp wait nobody /usr/sbin/tcpd in.tftpd /tftpboot
```

Debian packages will in general set this up correctly by default when they are installed.

> **Note:** Historically, TFTP servers used `/tftpboot` as directory to serve images from. However, Debian GNU/Linux packages may use other directories to comply with the Filesystem Hierarchy Standard (http://www.pathname.com/fhs/). For example, `tftpd-hpa` by default uses `/var/lib/tftpboot`. You may have to adjust the configuration examples in this section accordingly.

Look in `/etc/inetd.conf` and remember the directory which is used as the argument of **in.tftpd**[1]; you'll need that below. If you've had to change `/etc/inetd.conf`, you'll have to notify the running **inetd** process that the file has changed. On a Debian machine, run **`/etc/init.d/inetd reload`**; on other machines, find out the process ID for **inetd**, and run **`kill -HUP`** ***`inetd-pid`***.

4.3.4. Move TFTP Images Into Place

Next, place the TFTP boot image you need, as found in Section 4.2.1, in the **tftpd** boot image directory. You may have to make a link from that file to the file which **tftpd** will use for booting a particular client. Unfortunately, the file name is determined by the TFTP client, and there are no strong standards.

4.4. Automatic Installation

For installing on multiple computers it's possible to do fully automatic installations. Debian packages intended for this include `fai` (which uses an install server), `replicator`, `systemimager`, `autoinstall`, and the Debian Installer itself.

4.4.1. Automatic Installation Using the Debian Installer

The Debian Installer supports automating installs via preconfiguration files. A preconfiguration file can be loaded from the network or from removable media, and used to fill in answers to questions asked during the installation process.

Full documentation on preseeding including a working example that you can edit is in Appendix B.

1. The **`-l`** argument enables some versions of **in.tftpd** to log all requests to the system logs; this is useful for diagnosing boot errors.

Chapter 5. Booting the Installation System

5.1. Boot Parameters

Boot parameters are Linux kernel parameters which are generally used to make sure that peripherals are dealt with properly. For the most part, the kernel can auto-detect information about your peripherals. However, in some cases you'll have to help the kernel a bit.

If this is the first time you're booting the system, try the default boot parameters (i.e., don't try setting parameters) and see if it works correctly. It probably will. If not, you can reboot later and look for any special parameters that inform the system about your hardware.

Information on many boot parameters can be found in the Linux BootPrompt HOWTO (http://www.tldp.org/HOWTO/BootPrompt-HOWTO.html), including tips for obscure hardware. This section contains only a sketch of the most salient parameters. Some common gotchas are included below in Section 5.2.

When the kernel boots, a message

```
Memory:availk/totalk available
```

should be emitted early in the process. *total* should match the total amount of RAM, in kilobytes. If this doesn't match the actual amount of RAM you have installed, you need to use the **mem=*ram*** parameter, where *ram* is set to the amount of memory, suffixed with "k" for kilobytes, or "m" for megabytes. For example, both **mem=65536k** and **mem=64m** mean 64MB of RAM.

If you are booting with a serial console, generally the kernel will autodetect this. If you have a videocard (framebuffer) and a keyboard also attached to the computer which you wish to boot via serial console, you may have to pass the **console=*device*** argument to the kernel, where *device* is your serial device, which is usually something like `ttyS0`.

5.1.1. Debian Installer Parameters

The installation system recognizes a few additional boot parameters[1] which may be useful.

A number of parameters have a "short form" that helps avoid the limitations of the kernel command line options and makes entering the parameters easier. If a parameter has a short form, it will be listed in brackets behind the (normal) long form. Examples in this manual will normally use the short form too.

debconf/priority (priority)

This parameter sets the lowest priority of messages to be displayed.

The default installation uses **priority=high**. This means that both high and critical priority messages are shown, but medium and low priority messages are skipped. If problems are encountered, the installer adjusts the priority as needed.

If you add **priority=medium** as boot parameter, you will be shown the installation menu and gain more control over the installation. When **priority=low** is used, all messages are shown (this is equivalent to the *expert* boot method). With **priority=critical**, the installation system will display only critical messages and try to do the right thing without fuss.

1. With current kernels (2.6.9 or newer) you can use 32 command line options and 32 environment options. If these numbers are exceeded, the kernel will panic.

DEBIAN_FRONTEND

This boot parameter controls the type of user interface used for the installer. The current possible parameter settings are:

- **DEBIAN_FRONTEND=noninteractive**
- **DEBIAN_FRONTEND=text**
- **DEBIAN_FRONTEND=newt**
- **DEBIAN_FRONTEND=gtk**

The default frontend is **DEBIAN_FRONTEND=newt**. **DEBIAN_FRONTEND=text** may be preferable for serial console installs. Generally, only the **newt** frontend is available on default install media. On architectures that support it, the graphical installer uses the **gtk** frontend.

BOOT_DEBUG

Setting this boot parameter to 2 will cause the installer's boot process to be verbosely logged. Setting it to 3 makes debug shells available at strategic points in the boot process. (Exit the shells to continue the boot process.)

BOOT_DEBUG=0

This is the default.

BOOT_DEBUG=1

More verbose than usual.

BOOT_DEBUG=2

Lots of debugging information.

BOOT_DEBUG=3

Shells are run at various points in the boot process to allow detailed debugging. Exit the shell to continue the boot.

INSTALL_MEDIA_DEV

The value of the parameter is the path to the device to load the Debian installer from. For example, **INSTALL_MEDIA_DEV=/dev/floppy/0**

The boot floppy, which normally scans all floppies it can to find the root floppy, can be overridden by this parameter to only look at the one device.

debian-installer/framebuffer (fb)

Some architectures use the kernel framebuffer to offer installation in a number of languages. If framebuffer causes a problem on your system you can disable the feature by the parameter **fb=false**. Problem symptoms are error messages about bterm or bogl, a blank screen, or a freeze within a few minutes after starting the install.

Such problems have been reported on hppa.

debian-installer/theme (theme)

A theme determines how the user interface of the installer looks (colors, icons, etc.). What themes are available differs per frontend. Currently both the newt and gtk frontends only have a "dark" theme that was designed for visually impaired users. Set the theme by booting with **theme=*dark***.

debian-installer/probe/usb

Set to **`false`** to prevent probing for USB on boot, if that causes problems.

netcfg/disable_dhcp

By default, the `debian-installer` automatically probes for network configuration via DHCP. If the probe succeeds, you won't have a chance to review and change the obtained settings. You can get to the manual network setup only in case the DHCP probe fails.

If you have a DHCP server on your local network, but want to avoid it because e.g. it gives wrong answers, you can use the parameter **`netcfg/disable_dhcp=true`** to prevent configuring the network with DHCP and to enter the information manually.

hw-detect/start_pcmcia

Set to **`false`** to prevent starting PCMCIA services, if that causes problems. Some laptops are well known for this misbehavior.

preseed/url (url)

Specify the url to a preconfiguration file to download and use for automating the install. See Section 4.4.

preseed/file (file)

Specify the path to a preconfiguration file to load for automating the install. See Section 4.4.

preseed/interactive

Set to **`true`** to display questions even if they have been preseeded. Can be useful for testing or debugging a preconfiguration file. Note that this will have no effect on parameters that are passed as boot parameters, but for those a special syntax can be used. See Section B.5.2 for details.

auto-install/enable (auto)

Delay questions that are normally asked before preseeding is possible until after the network is configured. See Section B.2.3 for details about using this to automate installs.

cdrom-detect/eject

By default, before rebooting, `debian-installer` automatically ejects the optical media used during the installation. This can be unnecessary if the system does not automatically boot off the CD. In some cases it may even be undesirable, for example if the optical drive cannot reinsert the media itself and the user is not there to do it manually. Many slot loading, slim-line, and caddy style drives cannot reload media automatically.

Set to **`false`** to disable automatic ejection, and be aware that you may need to ensure that the system does not automatically boot from the optical drive after the initial installation.

debian-installer/allow_unauthenticated

By default the installer requires that repositories be authenticated using a known gpg key. Set to **`true`** to disable that authentication. **Warning: insecure, not recommended.**

rescue/enable

Set to **`true`** to enter rescue mode rather than performing a normal installation. See Section 8.7.

5.1.1.1. Using boot parameters to answer questions

With some exceptions, a value can be set at the boot prompt for any question asked during the installation, though this is only really useful in specific cases. General instructions how to do this can be found in Section B.2.2. Some specific examples are listed below.

debian-installer/locale (locale)

Can be used to set both the language and country for the installation. This will only work if the locale is supported in Debian. For example, use **`locale=de_CH`** to select German as language and Switzerland as country.

anna/choose_modules (modules)

Can be used to automatically load installer components that are not loaded by default. Examples of optional components that may be useful are `openssh-client-udeb` (so you can use **scp** during the installation) and `ppp-udeb` (see Section D.4).

netcfg/disable_dhcp

Set to **`true`** if you want to disable DHCP and instead force static network configuration.

mirror/protocol (protocol)

By default the installer will use the http protocol to download files from Debian mirrors and changing that to ftp is not possible during installations at normal priority. By setting this parameter to **`ftp`**, you can force the installer to use that protocol instead. Note that you cannot select an ftp mirror from a list, you have to enter the hostname manually.

tasksel:tasksel/first (tasks)

Can be used to select tasks that are not available from the interactive task list, such as the `kde-desktop` task. See Section 6.3.5.2 for additional information.

5.1.1.2. Passing parameters to kernel modules

If drivers are compiled into the kernel, you can pass parameters to them as described in the kernel documentation. However, if drivers are compiled as modules and because kernel modules are loaded a bit differently during an installation than when booting an installed system, it is not possible to pass parameters to modules as you would normally do. Instead, you need to use a special syntax recognized by the installer which will then make sure that the parameters are saved in the proper configuration files and will thus be used when the modules are actually loaded. The parameters will also be propagated automatically to the configuration for the installed system.

Note that it is now quite rare that parameters need to be passed to modules. In most cases the kernel will be able to probe the hardware present in a system and set good defaults that way. However, in some situations it may still be needed to set parameters manually.

The syntax to use to set parameters for modules is:

```
module_name.parameter_name=value
```

If you need to pass multiple parameters to the same or different modules, just repeat this. For example, to set an old 3Com network interface card to use the BNC (coax) connector and IRQ 10, you would pass:

```
3c509.xcvr=3 3c509.irq=10
```

5.1.1.3. Blacklisting kernel modules

Sometimes it may be necessary to blacklist a module to prevent it from being loaded automatically by the kernel and udev. One reason could be that a particular module causes problems with your hardware. The kernel also sometimes lists two different drivers for the same device. This can cause the device to not work correctly if the drivers conflict or if the wrong driver is loaded first.

You can blacklist a module using the following syntax: ***module_name*.blacklist=yes**. This will cause the module to be blacklisted in `/etc/modprobe.d/blacklist.local` both during the installation and for the installed system.

Note that a module may still be loaded by the installation system itself. You can prevent that from happening by running the installation in expert mode and unselecting the module from the list of modules displayed during the hardware detection phases.

5.2. Troubleshooting the Installation Process

5.2.1. CD-ROM Reliability

Sometimes, especially with older CD-ROM drives, the installer may fail to boot from a CD-ROM. The installer may also — even after booting successfully from CD-ROM — fail to recognize the CD-ROM or return errors while reading from it during the installation.

There are a many different possible causes for these problems. We can only list some common issues and provide general suggestions on how to deal with them. The rest is up to you.

There are two very simple things that you should try first.

- If the CD-ROM does not boot, check that it was inserted correctly and that it is not dirty.
- If the installer fails to recognize a CD-ROM, try just running the option Detect and mount CD-ROM a second time. Some DMA related issues with older CD-ROM drives are known to be resolved in this way.

If this does not work, then try the suggestions in the subsections below. Most, but not all, suggestions discussed there are valid for both CD-ROM and DVD, but we'll use the term CD-ROM for simplicity.

If you cannot get the installation working from CD-ROM, try one of the other installation methods that are available.

5.2.1.1. Common issues

- Some older CD-ROM drives do not support reading from discs that were burned at high speeds using a modern CD writer.

- If your system boots correctly from the CD-ROM, it does not necessarily mean that Linux also supports the CD-ROM (or, more correctly, the controller that your CD-ROM drive is connected to).
- Some older CD-ROM drives do not work correctly if “direct memory access” (DMA) is enabled.

5.2.1.2. How to investigate and maybe solve issues

If the CD-ROM fails to boot, try the suggestions listed below.

- Check that your BIOS actually supports booting from CD-ROM (older systems possibly don’t) and that your CD-ROM drive supports the media you are using.
- If you downloaded an iso image, check that the md5sum of that image matches the one listed for the image in the `MD5SUMS` file that should be present in the same location as where you downloaded the image from.

  ```
  $ md5sum debian-testing-i386-netinst.iso
  a20391b12f7ff22ef705cee4059c6b92  debian-testing-i386-netinst.iso
  ```

  Next, check that the md5sum of the burned CD-ROM matches as well. The following command should work. It uses the size of the image to read the correct number of bytes from the CD-ROM.

  ```
  $ dd if=/dev/cdrom | \
  > head -c ‘stat --format=%s debian-testing-i386-netinst.iso‘ | \
  > md5sum
  a20391b12f7ff22ef705cee4059c6b92  -
  262668+0 records in
  262668+0 records out
  134486016 bytes (134 MB) copied, 97.474 seconds, 1.4 MB/s
  ```

If, after the installer has been booted successfully, the CD-ROM is not detected, sometimes simply trying again may solve the problem. If you have more than one CD-ROM drive, try changing the CD-ROM to the other drive. If that does not work or if the CD-ROM is recognized but there are errors when reading from it, try the suggestions listed below. Some basic knowledge of Linux is required for this. To execute any of the commands, you should first switch to the second virtual console (VT2) and activate the shell there.

- Switch to VT4 or view the contents of `/var/log/syslog` (use **nano** as editor) to check for any specific error messages. After that, also check the output of **dmesg**.
- Check in the output of **dmesg** if your CD-ROM drive was recognized. You should see something like (the lines do not necessarily have to be consecutive):

  ```
  Probing IDE interface ide1...
  hdc: TOSHIBA DVD-ROM SD-R6112, ATAPI CD/DVD-ROM drive
  ide1 at 0x170-0x177,0x376 on irq 15
  hdc: ATAPI 24X DVD-ROM DVD-R CD-R/RW drive, 2048kB Cache, UDMA(33)
  Uniform CD-ROM driver Revision: 3.20
  ```

  If you don’t see something like that, chances are the controller your CD-ROM is connected to was not recognized or may be not supported at all. If you know what driver is needed for the controller, you can try loading it manually using **modprobe**.

- Check that there is a device node for your CD-ROM drive under `/dev/`. In the example above, this would be `/dev/hdc`. There should also be a `/dev/cdroms/cdrom0`.
- Use the **mount** command to check if the CD-ROM is already mounted; if not, try mounting it manually:

  ```
  $ mount /dev/hdc /cdrom
  ```

  Check if there are any error messages after that command.
- Check if DMA is currently enabled:

  ```
  $ cd /proc/ide/hdc
  $ grep using_dma settings
  using_dma      1       0       1       rw
  ```

  A "1" in the first column after `using_dma` means it is enabled. If it is, try disabling it:

  ```
  $ echo -n "using_dma:0" >settings
  ```

  Make sure that you are in the directory for the device that corresponds to your CD-ROM drive.
- If there are any problems during the installation, try checking the integrity of the CD-ROM using the option near the bottom of the installer's main menu. This option can also be used as a general test if the CD-ROM can be read reliably.

5.2.2. Boot Configuration

If you have problems and the kernel hangs during the boot process, doesn't recognize peripherals you actually have, or drives are not recognized properly, the first thing to check is the boot parameters, as discussed in Section 5.1.

Often, problems can be solved by removing add-ons and peripherals, and then trying booting again.

If you have a large amount of memory installed in your machine, more than 512M, and the installer hangs when booting the kernel, you may need to include a boot argument to limit the amount of memory the kernel sees, such as **`mem=512m`**.

5.2.3. Interpreting the Kernel Startup Messages

During the boot sequence, you may see many messages in the form `can't find` *`something`*, or *`something`* `not present`, `can't initialize` *`something`*, or even `this driver release depends on` *`something`*. Most of these messages are harmless. You see them because the kernel for the installation system is built to run on computers with many different peripheral devices. Obviously, no one computer will have every possible peripheral device, so the operating system may emit a few complaints while it looks for peripherals you don't own. You may also see the system pause for a while. This happens when it is waiting for a device to respond, and that device is not present on your system. If you find the time it takes to boot the system unacceptably long, you can create a custom kernel later (see Section 8.6).

5.2.4. Reporting Installation Problems

If you get through the initial boot phase but cannot complete the install, the menu option Save debug logs may be helpful. It lets you store system error logs and configuration information from the installer to a floppy, or download them using a web browser. This information may provide clues as to what went wrong and how to fix it. If you are submitting a bug report, you may want to attach this information to the bug report.

Other pertinent installation messages may be found in `/var/log/` during the installation, and `/var/log/installer/` after the computer has been booted into the installed system.

5.2.5. Submitting Installation Reports

If you still have problems, please submit an installation report. We also encourage installation reports to be sent even if the installation is successful, so that we can get as much information as possible on the largest number of hardware configurations.

Note that your installation report will be published in the Debian Bug Tracking System (BTS) and forwarded to a public mailing list. Make sure that you use an e-mail address that you do not mind being made public.

If you have a working Debian system, the easiest way to send an installation report is to install the `installation-report` and `reportbug` packages (**aptitude install installation-report reportbug**), configure `reportbug` as explained in Section 8.5.2, and run the command **reportbug installation-reports**.

Alternatively you can use this template when filling out installation reports, and file the report as a bug report against the `installation-reports` pseudo package, by sending it to `<submit@bugs.debian.org>`.

```
Package: installation-reports

Boot method: <How did you boot the installer? CD? floppy? network?>
Image version: <Full URL to image you downloaded is best>
Date: <Date and time of the install>

Machine: <Description of machine (eg, IBM Thinkpad R32)>
Processor:
Memory:
Partitions: <df -Tl will do; the raw partition table is preferred>

Output of lspci -nn and lspci -vnn:

Base System Installation Checklist:
[O] = OK, [E] = Error (please elaborate below), [ ] = didn't try it

Initial boot:           [ ]
Detect network card:    [ ]
Configure network:      [ ]
Detect CD:              [ ]
Load installer modules: [ ]
Detect hard drives:     [ ]
Partition hard drives:  [ ]
Install base system:    [ ]
Clock/timezone setup:   [ ]
User/password setup:    [ ]
```

```
Install tasks:          [ ]
Install boot loader:    [ ]
Overall install:        [ ]

Comments/Problems:

<Description of the install, in prose, and any thoughts, comments
      and ideas you had during the initial install.>
```

In the bug report, describe what the problem is, including the last visible kernel messages in the event of a kernel hang. Describe the steps that you did which brought the system into the problem state.

Chapter 6. Using the Debian Installer

6.1. How the Installer Works

The Debian Installer consists of a number of special-purpose components to perform each installation task. Each component performs its task, asking the user questions as necessary to do its job. The questions themselves are given priorities, and the priority of questions to be asked is set when the installer is started.

When a default installation is performed, only essential (high priority) questions will be asked. This results in a highly automated installation process with little user interaction. Components are automatically run in sequence; which components are run depends mainly on the installation method you use and on your hardware. The installer will use default values for questions that are not asked.

If there is a problem, the user will see an error screen, and the installer menu may be shown in order to select some alternative action. If there are no problems, the user will never see the installer menu, but will simply answer questions for each component in turn. Serious error notifications are set to priority "critical" so the user will always be notified.

Some of the defaults that the installer uses can be influenced by passing boot arguments when `debian-installer` is started. If, for example, you wish to force static network configuration (DHCP is used by default if available), you could add the boot parameter **`netcfg/disable_dhcp=true`**. See Section 5.1.1 for available options.

Power users may be more comfortable with a menu-driven interface, where each step is controlled by the user rather than the installer performing each step automatically in sequence. To use the installer in a manual, menu-driven way, add the boot argument **`priority=medium`**.

If your hardware requires you to pass options to kernel modules as they are installed, you will need to start the installer in "expert" mode. This can be done by either using the **expert** command to start the installer or by adding the boot argument **`priority=low`**. Expert mode gives you full control over `debian-installer`.

The normal installer display is character-based (as opposed to the now more familiar graphical interface). The mouse is not operational in this environment. Here are the keys you can use to navigate within the various dialogs. The **Tab** or **right** arrow keys move "forward", and the **Shift-Tab** or **left** arrow keys move "backward" between displayed buttons and selections. The **up** and **down** arrow select different items within a scrollable list, and also scroll the list itself. In addition, in long lists, you can type a letter to cause the list to scroll directly to the section with items starting with the letter you typed and use **Pg-Up** and **Pg-Down** to scroll the list in sections. The **space bar** selects an item such as a checkbox. Use **Enter** to activate choices.

Error messages and logs are redirected to the fourth console. You can access this console by pressing **Left Alt-F4** (hold the left **Alt** key while pressing the **F4** function key); get back to the main installer process with **Left Alt-F1**.

These messages can also be found in `/var/log/syslog`. After installation, this log is copied to `/var/log/installer/syslog` on your new system. Other installation messages may be found in `/var/log/` during the installation, and `/var/log/installer/` after the computer has been booted into the installed system.

6.2. Components Introduction

Here is a list of installer components with a brief description of each component's purpose. Details you might need to know about using a particular component are in Section 6.3.

main-menu

Shows the list of components to the user during installer operation, and starts a component when it is selected. Main-menu's questions are set to priority medium, so if your priority is set to high or critical (high is the default), you will not see the menu. On the other hand, if there is an error which requires your intervention, the question priority may be downgraded temporarily to allow you to resolve the problem, and in that case the menu may appear.

You can get to the main menu by selecting the "Back" button repeatedly to back all the way out of the currently running component.

localechooser

Allows the user to select localization options for the installation and the installed system: language, country and locales. The installer will display messages in the selected language, unless the translation for that language is not complete in which case some messages may be shown in English.

kbd-chooser

Shows a list of keyboards, from which the user chooses the model which matches his own.

hw-detect

Automatically detects most of the system's hardware, including network cards, disk drives, and PCMCIA.

cdrom-detect

Looks for and mounts a Debian installation CD.

netcfg

Configures the computer's network connections so it can communicate over the internet.

iso-scan

Looks for ISO file systems, which may be on a CD-ROM or on the hard drive.

choose-mirror

Presents a list of Debian archive mirrors. The user may choose the source of his installation packages.

cdrom-checker

Checks integrity of a CD-ROM. This way, the user may assure him/herself that the installation CD-ROM was not corrupted.

lowmem

Lowmem tries to detect systems with low memory and then does various tricks to remove unnecessary parts of `debian-installer` from the memory (at the cost of some features).

anna

Anna's Not Nearly APT. Installs packages which have been retrieved from the chosen mirror or CD.

partman

Allows the user to partition disks attached to the system, create file systems on the selected partitions, and attach them to the mountpoints. Included are also interesting features like a fully automatic mode or LVM support. This is the preferred partitioning tool in Debian.

autopartkit

Automatically partitions an entire disk according to preset user preferences.

partitioner

Allows the user to partition disks attached to the system. A partitioning program appropriate to your computer's architecture is chosen.

partconf

Displays a list of partitions, and creates file systems on the selected partitions according to user instructions.

lvmcfg

Helps the user with the configuration of the *LVM* (Logical Volume Manager).

mdcfg

Allows the user to setup Software *RAID* (Redundant Array of Inexpensive Disks). This Software RAID is usually superior to the cheap IDE (pseudo hardware) RAID controllers found on newer motherboards.

tzsetup

Selects the time zone, based on the location selected earlier.

clock-setup

Determines whether the clock is set to UTC or not.

user-setup

Sets up the root password, and adds a non-root user.

base-installer

Installs the most basic set of packages which would allow the computer to operate under Linux when rebooted.

apt-setup

Configures apt, mostly automatically, based on what media the installer is running from.

pkgsel

Uses `tasksel` to select and install additional software.

os-prober

Detects currently installed operating systems on the computer and passes this information to the bootloader-installer, which may offer you an ability to add discovered operating systems to the bootloader's start menu. This way the user could easily choose at the boot time which operating system to start.

bootloader-installer

The various bootloader installers each install a boot loader program on the hard disk, which is necessary for the computer to start up using Linux without using a floppy or CD-ROM. Many boot loaders allow the user to choose an alternate operating system each time the computer boots.

shell

Allows the user to execute a shell from the menu, or in the second console.

save-logs

Provides a way for the user to record information on a floppy disk, network, hard disk, or other media when trouble is encountered, in order to accurately report installer software problems to Debian developers later.

6.3. Using Individual Components

In this section we will describe each installer component in detail. The components have been grouped into stages that should be recognizable for users. They are presented in the order they appear during the install. Note that not all modules will be used for every installation; which modules are actually used depends on the installation method you use and on your hardware.

6.3.1. Setting up Debian Installer and Hardware Configuration

Let's assume the Debian Installer has booted and you are facing its first screen. At this time, the capabilities of `debian-installer` are still quite limited. It doesn't know much about your hardware, preferred language, or even the task it should perform. Don't worry. Because `debian-installer` is quite clever, it can automatically probe your hardware, locate the rest of its components and upgrade itself to a capable installation system. However, you still need to help `debian-installer` with some information it can't determine automatically (like selecting your preferred language, keyboard layout or desired network mirror).

You will notice that `debian-installer` performs *hardware detection* several times during this stage. The first time is targeted specifically at the hardware needed to load installer components (e.g. your CD-ROM or network card). As not all drivers may be available during this first run, hardware detection needs to be repeated later in the process.

6.3.1.1. Check available memory / low memory mode

One of the first things `debian-installer` does, is to check available memory. If the available memory is limited, this component will make some changes in the installation process which hopefully will allow you to install Debian GNU/Linux on your system.

The first measure taken to reduce memory consumption by the installer is to disable translations, which means that the installation can only be done in English. Of course, you can still localize the installed system after the installation has completed.

If that is not sufficient, the installer will further reduce memory consumption by loading only those components essential to complete a basic installation. This reduces the functionality of the installation system. You will be given the opportunity to load additional components manually, but you should be aware that each component you select will use additional memory and thus may cause the installation to fail.

If the installer runs in low memory mode, it is recommended to create a relatively large swap partition (64–128MB). The swap partition will be used as virtual memory and thus increases the amount of memory available to the system. The installer will activate the swap partition as early as possible in the installation process. Note that heavy use of swap will reduce performance of your system and may lead to high disk activity.

Despite these measures, it is still possible that your system freezes, that unexpected errors occur or that processes are killed by the kernel because the system runs out of memory (which will result in "Out of memory" messages on VT4 and in the syslog).

For example, it has been reported that creating a big ext3 file system fails in low memory mode when there is insufficient swap space. If a larger swap doesn't help, try creating the file system as ext2 (which is an essential component of the installer) instead. It is possible to change an ext2 partition to ext3 after the installation.

6.3.1.2. Selecting Localization Options

In most cases the first questions you will be asked concern the selection of localization options to be used both for the installation and for the installed system. The localization options consist of language, country and locales.

The language you choose will be used for the rest of the installation process, provided a translation of the different dialogs is available. If no valid translation is available for the selected language, the installer will default to English.

The selected country will be used later in the installation process to pick the default timezone and a Debian mirror appropriate for your geographic location. Language and country together will be used to set the default locale for your system and to help select your keyboard.

You will first be asked to select your preferred language. The language names are listed in both English (left side) and in the language itself (right side); the names on the right side are also shown in the proper script for the language. The list is sorted on the English names. At the top of the list is an extra option that allows you to select the "C" locale instead of a language. Choosing the "C" locale will result in the installation proceding in English; the installed system will have no localization support as the `locales` package will not be installed.

If you selected a language that is recognized as an official language for more than one country[1], you will next be asked to select a country. If you choose Other at the bottom of the list, you will be presented with a list of all countries, grouped by continent. If the language has only one country associated with it, that country will be selected automatically.

A default locale will be selected based on the selected language and country. If you are installing at medium or low priority, you will have the option of selecting a different default locale and of selecting additional locales to be generated for the installed system.

1. In technical terms: where multiple locales exist for that language with differing country codes.

6.3.1.3. Choosing a Keyboard

Keyboards are often tailored to the characters used in a language. Select a layout that conforms to the keyboard you are using, or select something close if the keyboard layout you want isn't represented. Once the system installation is complete, you'll be able to select a keyboard layout from a wider range of choices (run **kbdconfig** as root after you have completed the installation).

Move the highlight to the keyboard selection you desire and press **Enter**. Use the arrow keys to move the highlight — they are in the same place in all national language keyboard layouts, so they are independent of the keyboard configuration. An 'extended' keyboard is one with **F1** through **F10** keys along the top row.

6.3.1.4. Looking for the Debian Installer ISO Image

When installing via the *hd-media* method, there will be a moment where you need to find and mount the Debian Installer iso image in order to get the rest of the installation files. The component **iso-scan** does exactly this.

At first, **iso-scan** automatically mounts all block devices (e.g. partitions) which have some known filesystem on them and sequentially searches for filenames ending with `.iso` (or `.ISO` for that matter). Beware that the first attempt scans only files in the root directory and in the first level of subdirectories (i.e. it finds `/`*`whatever`*`.iso`, `/data/`*`whatever`*`.iso`, but not `/data/tmp/`*`whatever`*`.iso`). After an iso image has been found, **iso-scan** checks its content to determine if the image is a valid Debian iso image or not. In the former case we are done, in the latter **iso-scan** seeks for another image.

In case the previous attempt to find an installer iso image fails, **iso-scan** will ask you whether you would like to perform a more thorough search. This pass doesn't just look into the topmost directories, but really traverses whole filesystem.

If **iso-scan** does not discover your installer iso image, reboot back to your original operating system and check if the image is named correctly (ending in `.iso`), if it is placed on a filesystem recognizable by `debian-installer`, and if it is not corrupted (verify the checksum). Experienced Unix users could do this without rebooting on the second console.

6.3.1.5. Configuring Network

As you enter this step, if the system detects that you have more than one network device, you'll be asked to choose which device will be your *primary* network interface, i.e. the one which you want to use for installation. The other interfaces won't be configured at this time. You may configure additional interfaces after installation is complete; see the interfaces(5) man page.

By default, `debian-installer` tries to configure your computer's network automatically via DHCP. If the DHCP probe succeeds, you are done. If the probe fails, it may be caused by many factors ranging from unplugged network cable, to a misconfigured DHCP setup. Or maybe you don't have a DHCP server in your local network at all. For further explanation, check the error messages on the fourth console. In any case, you will be asked if you want to retry, or if you want to perform a manual setup. DHCP servers are sometimes really slow in their responses, so if you are sure everything is in place, try again.

The manual network setup in turn asks you a number of questions about your network, notably `IP address`, `Netmask`, `Gateway`, `Name server addresses`, and a `Hostname`. Moreover, if you have a wireless network interface, you will be asked to provide your `Wireless ESSID` and a `WEP key`. Fill in the answers from Section 3.3.

Note: Some technical details you might, or might not, find handy: the program assumes the network IP address is the bitwise-AND of your system's IP address and your netmask. It will guess the broadcast address is the bitwise OR of your system's IP address with the bitwise negation of the netmask. It will also guess your gateway. If you can't find any of these answers, use the system's guesses — you can change them once the system has been installed, if necessary, by editing `/etc/network/interfaces`.

6.3.2. Partitioning and Mount Point Selection

At this time, after hardware detection has been executed a final time, `debian-installer` should be at its full strength, customized for the user's needs and ready to do some real work. As the title of this section indicates, the main task of the next few components lies in partitioning your disks, creating filesystems, assigning mountpoints and optionally configuring closely related issues like LVM or RAID devices.

6.3.2.1. Partitioning Your Disks

Now it is time to partition your disks. If you are uncomfortable with partitioning, or just want to know more details, see Appendix C.

First you will be given the opportunity to automatically partition either an entire drive, or available free space on a drive. This is also called "guided" partitioning. If you do not want to autopartition, choose Manual from the menu.

If you choose guided partitioning, you may have three options: to create partitions directly on the hard disk (classic method), or to create them using Logical Volume Management (LVM), or to create them using encrypted LVM[2].

Note: The option to use (encrypted) LVM may not be available on all architectures.

When using LVM or encrypted LVM, the installer will create most partitions inside one big partition; the advantage of this method is that partitions inside this big partition can be resized relatively easily later. In the case of encrypted LVM the big partition will not be readable without knowing a special key phrase, thus providing extra security of your (personal) data.

When using encrypted LVM, the installer will also automatically erase the disk by writing random data to it. This further improves security (as it makes it impossible to tell which parts of the disk are in use and also makes sure that any traces of previous installations are erased), but may take some time depending on the size of your disk.

Note: If you choose guided partitioning using LVM or encrypted LVM, some changes in the partition table will need to be written to the selected disk while LVM is being set up. These changes effectively erase all data that is currently on the selected hard disk and you will not be able to undo them later. However, the installer will ask you to confirm these changes before they are written to disk.

2. The installer will encrypt the LVM volume group using a 256 bit AES key and makes use of the kernel's "dm-crypt" support.

If you choose guided partitioning (either classic or using (encrypted) LVM) for a whole disk, you will first be asked to select the disk you want to use. Check that all your disks are listed and, if you have several disks, make sure you select the correct one. The order they are listed in may differ from what you are used to. The size of the disks may help to identify them.

Any data on the disk you select will eventually be lost, but you will always be asked to confirm any changes before they are written to the disk. If you have selected the classic method of partitioning, you will be able to undo any changes right until the end; when using (encrypted) LVM this is not possible.

Next, you will be able to choose from the schemes listed in the table below. All schemes have their pros and cons, some of which are discussed in Appendix C. If you are unsure, choose the first one. Bear in mind that guided partitioning needs a certain minimal amount of free space to operate with. If you don't give it at least about 1GB of space (depends on chosen scheme), guided partitioning will fail.

Partitioning scheme	Minimum space	Created partitions
All files in one partition	600MB	`/`, swap
Separate /home partition	500MB	`/`, `/home`, swap
Separate /home, /usr, /var and /tmp partitions	1GB	`/`, `/home`, `/usr`, `/var`, `/tmp`, swap

If you choose guided partitioning using (encrypted) LVM, the installer will also create a separate `/boot` partition. The other partitions, including the swap partition, will be created inside the LVM partition.

After selecting a scheme, the next screen will show your new partition table, including information on whether and how partitions will be formatted and where they will be mounted.

The list of partitions might look like this:

```
IDE1 master (hda) - 6.4 GB WDC AC36400L
      #1 primary   16.4 MB  B f ext2       /boot
      #2 primary  551.0 MB      swap       swap
      #3 primary    5.8 GB      ntfs
         pri/log    8.2 MB      FREE SPACE

IDE1 slave (hdb) - 80.0 GB ST380021A
      #1 primary   15.9 MB      ext3
      #2 primary  996.0 MB      fat16
      #3 primary    3.9 GB      xfs        /home
      #5 logical    6.0 GB    f ext3       /
      #6 logical    1.0 GB    f ext3       /var
      #7 logical  498.8 MB      ext3
      #8 logical  551.5 MB      swap       swap
      #9 logical   65.8 GB      ext2
```

This example shows two IDE harddrives divided into several partitions; the first disk has some free space. Each partition line consists of the partition number, its type, size, optional flags, file system, and mountpoint (if any). Note: this particular setup cannot be created using guided partitioning but it does show possible variation that can be achieved using manual partitioning.

This concludes the guided partitioning. If you are satisfied with the generated partition table, you can choose Finish partitioning and write changes to disk from the menu to implement the new partition table (as described at the end of this section). If you are not happy, you can choose to Undo changes

to partitions and run guided partitioning again, or modify the proposed changes as described below for manual partitioning.

A similar screen to the one shown just above will be displayed if you choose manual partitioning except that your existing partition table will be shown and without the mount points. How to manually setup your partition table and the usage of partitions by your new Debian system will be covered in the remainder of this section.

If you select a pristine disk which has neither partitions nor free space on it, you will be asked if a new partition table should be created (this is needed so you can create new partitions). After this, a new line entitled "FREE SPACE" should appear in the table under the selected disk.

If you select some free space, you will have the opportunity to create a new partition. You will have to answer a quick series of questions about its size, type (primary or logical), and location (beginning or end of the free space). After this, you will be presented with a detailed overview of your new partition. The main setting is Use as:, which determines if the partition will have a file system on it, or be used for swap, software RAID, LVM, an encrypted file system, or not be used at all. Other settings include mountpoint, mount options, and bootable flag; which settings are shown depends on how the partition is to be used. If you don't like the preselected defaults, feel free to change them to your liking. E.g. by selecting the option Use as:, you can choose a different filesystem for this partition, including options to use the partition for swap, software RAID, LVM, or not use it at all. Another nice feature is the ability to copy data from an existing partition onto this one. When you are satisfied with your new partition, select Done setting up the partition and you will return to **partman**'s main screen.

If you decide you want to change something about your partition, simply select the partition, which will bring you to the partition configuration menu. This is the same screen as is used when creating a new partition, so you can change the same settings. One thing that may not be very obvious at a first glance is that you can resize the partition by selecting the item displaying the size of the partition. Filesystems known to work are at least fat16, fat32, ext2, ext3 and swap. This menu also allows you to delete a partition.

Be sure to create at least two partitions: one for the *root* filesystem (which must be mounted as `/`) and one for *swap*. If you forget to mount the root filesystem, **partman** won't let you continue until you correct this issue.

Capabilities of **partman** can be extended with installer modules, but are dependent on your system's architecture. So if you can't see all promised goodies, check if you have loaded all required modules (e.g. `partman-ext3`, `partman-xfs`, or `partman-lvm`).

After you are satisfied with partitioning, select Finish partitioning and write changes to disk from the partitioning menu. You will be presented with a summary of changes made to the disks and asked to confirm that the filesystems should be created as requested.

6.3.2.2. Configuring Multidisk Device (Software RAID)

If you have more than one harddrive[3] in your computer, you can use **mdcfg** to setup your drives for increased performance and/or better reliability of your data. The result is called *Multidisk Device* (or after its most famous variant *software RAID*).

MD is basically a bunch of partitions located on different disks and combined together to form a *logical* device. This device can then be used like an ordinary partition (i.e. in **partman** you can format it, assign a mountpoint, etc.).

3. To be honest, you can construct MD device even from partitions residing on single physical drive, but that won't bring you anything useful.

What benefits this brings depends on the type of MD device you are creating. Currently supported are:

RAID0

Is mainly aimed at performance. RAID0 splits all incoming data into *stripes* and distributes them equally over each disk in the array. This can increase the speed of read/write operations, but when one of the disks fails, you will lose *everything* (part of the information is still on the healthy disk(s), the other part *was* on the failed disk).

The typical use for RAID0 is a partition for video editing.

RAID1

Is suitable for setups where reliability is the first concern. It consists of several (usually two) equally-sized partitions where every partition contains exactly the same data. This essentially means three things. First, if one of your disks fails, you still have the data mirrored on the remaining disks. Second, you can use only a fraction of the available capacity (more precisely, it is the size of the smallest partition in the RAID). Third, file-reads are load-balanced among the disks, which can improve performance on a server, such as a file server, that tends to be loaded with more disk reads than writes.

Optionally you can have a spare disk in the array which will take the place of the failed disk in the case of failure.

RAID5

Is a good compromise between speed, reliability and data redundancy. RAID5 splits all incoming data into stripes and distributes them equally on all but one disk (similar to RAID0). Unlike RAID0, RAID5 also computes *parity* information, which gets written on the remaining disk. The parity disk is not static (that would be called RAID4), but is changing periodically, so the parity information is distributed equally on all disks. When one of the disks fails, the missing part of information can be computed from remaining data and its parity. RAID5 must consist of at least three active partitions. Optionally you can have a spare disk in the array which will take the place of the failed disk in the case of failure.

As you can see, RAID5 has a similar degree of reliability to RAID1 while achieving less redundancy. On the other hand, it might be a bit slower on write operations than RAID0 due to computation of parity information.

To sum it up:

Type	**Minimum Devices**	**Spare Device**	**Survives disk failure?**	**Available Space**
RAID0	2	no	no	Size of the smallest partition multiplied by number of devices in RAID
RAID1	2	optional	yes	Size of the smallest partition in RAID

Type	Minimum Devices	Spare Device	Survives disk failure?	Available Space
RAID5	3	optional	yes	Size of the smallest partition multiplied by (number of devices in RAID minus one)

If you want to know more about Software RAID, have a look at Software RAID HOWTO (http://www.tldp.org/HOWTO/Software-RAID-HOWTO.html).

To create a MD device, you need to have the desired partitions it should consist of marked for use in a RAID. (This is done in **partman** in the Partition settings menu where you should select Use as:⟶physical volume for RAID.)

Warning

Support for MD is a relatively new addition to the installer. You may experience problems for some RAID levels and in combination with some bootloaders if you try to use MD for the root (`/`) filesystem. For experienced users, it may be possible to work around some of these problems by executing some configuration or installation steps manually from a shell.

Next, you should choose Configure software RAID from the main **partman** menu. (The menu will only appear after you mark at least one partition for use as physical volume for RAID.) On the first screen of **mdcfg** simply select Create MD device. You will be presented with a list of supported types of MD devices, from which you should choose one (e.g. RAID1). What follows depends on the type of MD you selected.

- RAID0 is simple — you will be issued with the list of available RAID partitions and your only task is to select the partitions which will form the MD.
- RAID1 is a bit more tricky. First, you will be asked to enter the number of active devices and the number of spare devices which will form the MD. Next, you need to select from the list of available RAID partitions those that will be active and then those that will be spare. The count of selected partitions must be equal to the number provided few seconds ago. Don't worry. If you make a mistake and select different number of partitions, the `debian-installer` won't let you continue until you correct the issue.
- RAID5 has a similar setup procedure as RAID1 with the exception that you need to use at least *three* active partitions.

It is perfectly possible to have several types of MD at once. For example, if you have three 200 GB hard drives dedicated to MD, each containing two 100 GB partitions, you can combine the first partitions on all three disks into the RAID0 (fast 300 GB video editing partition) and use the other three partitions (2 active and 1 spare) for RAID1 (quite reliable 100 GB partition for `/home`).

After you setup MD devices to your liking, you can Finish **mdcfg** to return back to the **partman** to create filesystems on your new MD devices and assign them the usual attributes like mountpoints.

6.3.2.3. Configuring the Logical Volume Manager (LVM)

If you are working with computers at the level of system administrator or "advanced" user, you have surely seen the situation where some disk partition (usually the most important one) was short on space, while some other partition was grossly underused and you had to manage this situation with moving stuff around, symlinking, etc.

To avoid the described situation you can use Logical Volume Manager (LVM). Simply said, with LVM you can combine your partitions (*physical volumes* in LVM lingo) to form a virtual disk (so called *volume group*), which can then be divided into virtual partitions (*logical volumes*). The point is that logical volumes (and of course underlying volume groups) can span across several physical disks.

Now when you realize you need more space for your old 160GB `/home` partition, you can simply add a new 300GB disk to the computer, join it with your existing volume group and then resize the logical volume which holds your `/home` filesystem and voila — your users have some room again on their renewed 460GB partition. This example is of course a bit oversimplified. If you haven't read it yet, you should consult the LVM HOWTO (http://www.tldp.org/HOWTO/LVM-HOWTO.html).

LVM setup in `debian-installer` is quite simple and completely supported inside **partman**. First, you have to mark the partition(s) to be used as physical volumes for LVM. This is done in the Partition settings menu where you should select Use as:⟶physical volume for LVM.

When you return to the main **partman** screen, you will see a new option Configure the Logical Volume Manager. When you select that, you will first be asked to confirm pending changes to the partition table (if any) and after that the LVM configuration menu will be shown. Above the menu a summary of the LVM configuration is shown. The menu itself is context sensitive and only shows valid actions. The possible actions are:

- Display configuration details: shows LVM device structure, names and sizes of logical volumes and more
- Create volume group
- Create logical volume
- Delete volume group
- Delete logical volume
- Extend volume group
- Reduce volume group
- Finish: return to the main **partman** screen

Use the options in that menu to first create a volume group and then create your logical volumes inside it.

After you return to the main **partman** screen, any created logical volumes will be displayed in the same way as ordinary partitions (and you should treat them as such).

6.3.2.4. Configuring Encrypted Volumes

`debian-installer` allows you to set up encrypted partitions. Every file you write to such a partition is immediately saved to the device in encrypted form. Access to the encrypted data is granted only after entering the *passphrase* used when the encrypted partition was originally created. This feature is useful to protect sensitive data in case your laptop or hard drive gets stolen. The thief might get

physical access to the hard drive, but without knowing the right passphrase, the data on the hard drive will look like random characters.

The two most important partitions to encrypt are: the home partition, where your private data resides, and the swap partition, where sensitive data might be stored temporarily during operation. Of course, nothing prevents you from encrypting any other partitions that might be of interest. For example `/var` where database servers, mail servers or print servers store their data, or `/tmp` which is used by various programs to store potentially interesting temporary files. Some people may even want to encrypt their whole system. The only exception is the `/boot` partition which must remain unencrypted, because currently there is no way to load the kernel from an encrypted partition.

> **Note:** Please note that the performance of encrypted partitions will be less than that of unencrypted ones because the data needs to be decrypted or encrypted for every read or write. The performance impact depends on your CPU speed, chosen cipher and a key length.

To use encryption, you have to create a new partition by selecting some free space in the main partitioning menu. Another option is to choose an existing partition (e.g. a regular partition, an LVM logical volume or a RAID volume). In the Partition settings menu, you need to select physical volume for encryption at the Use as: option. The menu will then change to include several cryptographic options for the partition.

`debian-installer` supports several encryption methods. The default method is *dm-crypt* (included in newer Linux kernels, able to host LVM physical volumes), the other is *loop-AES* (older, maintained separately from the Linux kernel tree). Unless you have compelling reasons to do otherwise, it is recommended to use the default.

First, let's have a look at the options available when you select **`Device-mapper (dm-crypt)`** as the encryption method. As always: when in doubt, use the defaults, because they have been carefully chosen with security in mind.

Encryption: **`aes`**

This option lets you select the encryption algorithm (*cipher*) which will be used to encrypt the data on the partition. `debian-installer` currently supports the following block ciphers: *aes*, *blowfish*, *serpent*, and *twofish*. It is beyond the scope of this document to discuss the qualities of these different algorithms, however, it might help your decision to know that in 2000, *AES* was chosen by the American National Institute of Standards and Technology as the standard encryption algorithm for protecting sensitive information in the 21st century.

Key size: **`256`**

Here you can specify the length of the encryption key. With a larger key size, the strength of the encryption is generally improved. On the other hand, increasing the length of the key usually has a negative impact on performance. Available key sizes vary depending on the cipher.

IV algorithm: **`cbc-essiv:sha256`**

The *Initialization Vector* or *IV* algorithm is used in cryptography to ensure that applying the cipher on the same *clear text* data with the same key always produces a unique *cipher text*. The idea is to prevent the attacker from deducing information from repeated patterns in the encrypted data.

From the provided alternatives, the default **`cbc-essiv:sha256`** is currently the least vulnerable to known attacks. Use the other alternatives only when you need to ensure compatibility with some previously installed system that is not able to use newer algorithms.

Encryption key: **`Passphrase`**

Here you can choose the type of the encryption key for this partition.

Passphrase

The encryption key will be computed[4] on the basis of a passphrase which you will be able to enter later in the process.

Random key

A new encryption key will be generated from random data each time you try to bring up the encrypted partition. In other words: on every shutdown the content of the partition will be lost as the key is deleted from memory. (Of course, you could try to guess the key with a brute force attack, but unless there is an unknown weakness in the cipher algorithm, it is not achievable in our lifetime.)

Random keys are useful for swap partitions because you do not need to bother yourself with remembering the passphrase or wiping sensitive information from the swap partition before shutting down your computer. However, it also means that you will *not* be able to use the "suspend-to-disk" functionality offered by newer Linux kernels as it will be impossible (during a subsequent boot) to recover the suspended data written to the swap partition.

Erase data: **`yes`**

Determines whether the content of this partition should be overwritten with random data before setting up the encryption. This is recommended because it might otherwise be possible for an attacker to discern which parts of the partition are in use and which are not. In addition, this will make it harder to recover any leftover data from previous installations[5].

If you select Encryption method:⟶Loopback (loop-AES), the menu changes to provide the following options:

Encryption: **`AES256`**

For loop-AES, unlike dm-crypt, the options for cipher and key size are combined, so you can select both at the same time. Please see the above sections on ciphers and key sizes for further information.

Encryption key: **`Keyfile (GnuPG)`**

Here you can select the type of the encryption key for this partition.

Keyfile (GnuPG)

The encryption key will be generated from random data during the installation. Moreover this key will be encrypted with GnuPG, so to use it, you will need to enter the proper passphrase (you will be asked to provide one later in the process).

Random key

Please see the section on random keys above.

4. Using a passphrase as the key currently means that the partition will be set up using LUKS (http://luks.endorphin.org/).
5. It is believed that the guys from three-letter agencies can restore the data even after several rewrites of the magnetooptical media, though.

Erase data: **yes**

Please see the the section on erasing data above.

After you have selected the desired parameters for your encrypted partitions, return back to the main partitioning menu. There should now be a new menu item called Configure encrypted volumes. After you select it, you will be asked to confirm the deletion of data on partitions marked to be erased and possibly other actions such as writing a new partition table. For large partitions this might take some time.

Next you will be asked to enter a passphrase for partitions configured to use one. Good passphrases should be longer than 8 characters, should be a mixture of letters, numbers and other characters and should not contain common dictionary words or information easily associable with you (such as birthdates, hobbies, pet names, names of family members or relatives, etc.).

Warning

Before you input any passphrases, you should have made sure that your keyboard is configured correctly and generates the expected characters. If you are unsure, you can switch to the second virtual console and type some text at the prompt. This ensures that you won't be surprised later, e.g. by trying to input a passphrase using a qwerty keyboard layout when you used an azerty layout during the installation. This situation can have several causes. Maybe you switched to another keyboard layout during the installation, or the selected keyboard layout might not have been set up yet when entering the passphrase for the root file system.

If you selected to use methods other than a passphrase to create encryption keys, they will be generated now. Because the kernel may not have gathered a sufficient amount of entropy at this early stage of the installation, the process may take a long time. You can help speed up the process by generating entropy: e.g. by pressing random keys, or by switching to the shell on the second virtual console and generating some network and disk traffic (downloading some files, feeding big files into `/dev/null`, etc.). This will be repeated for each partition to be encrypted.

After returning to the main partitioning menu, you will see all encrypted volumes as additional partitions which can be configured in the same way as ordinary partitions. The following example shows two different volumes. The first one is encrypted via dm-crypt, the second one via loop-AES.

```
Encrypted volume (sda2_crypt) - 115.1 GB Linux device-mapper
     #1 115.1 GB  F ext3

Loopback (loop0) - 515.2 MB AES256 keyfile
     #1 515.2 MB  F ext3
```

Now is the time to assign mount points to the volumes and optionally change the file system types if the defaults do not suit you.

One thing to note here are the identifiers in parentheses (*`sda2_crypt`* and *`loop0`* in this case) and the mount points you assigned to each encrypted volume. You will need this information later when booting the new system. The differences between ordinary boot process and boot process with encryption involved will be covered later in Section 7.2.

Once you are satisfied with the partitioning scheme, continue with the installation.

6.3.3. Setting up the System

After partitioning the installer asks a few more questions that will be used to set up the system it is about to install.

6.3.3.1. Configuring Your Time Zone

Depending on the location selected at the beginning of the installation process, you might be shown a list of timezones relevant for that location. If your location has only one time zone, you will not be asked anything and the system will assume that time zone.

6.3.3.2. Configuring the Clock

The installer might ask you if the computer's clock is set to UTC. Normally this question is avoided if possible and the installer tries to work out whether the clock is set to UTC based on things like what other operating systems are installed.

In expert mode you will always be able to choose whether or not the clock is set to UTC.

Note that the installer does not currently allow you to actually set the time in the computer's clock. You can set the clock to the current time after you have installed, if it is incorrect or if it was previously not set to UTC.

6.3.3.3. Setting Up Users And Passwords

6.3.3.3.1. Set the Root Password

The *root* account is also called the *super-user*; it is a login that bypasses all security protection on your system. The root account should only be used to perform system administration, and only used for as short a time as possible.

Any password you create should contain at least 6 characters, and should contain both upper- and lower-case characters, as well as punctuation characters. Take extra care when setting your root password, since it is such a powerful account. Avoid dictionary words or use of any personal information which could be guessed.

If anyone ever tells you they need your root password, be extremely wary. You should normally never give your root password out, unless you are administering a machine with more than one system administrator.

6.3.3.3.2. Create an Ordinary User

The system will ask you whether you wish to create an ordinary user account at this point. This account should be your main personal log-in. You should *not* use the root account for daily use or as your personal login.

Why not? Well, one reason to avoid using root's privileges is that it is very easy to do irreparable damage as root. Another reason is that you might be tricked into running a *Trojan-horse* program — that is a program that takes advantage of your super-user powers to compromise the security of your system behind your back. Any good book on Unix system administration will cover this topic in more detail — consider reading one if it is new to you.

You will first be prompted for the user's full name. Then you'll be asked for a name for the user account; generally your first name or something similar will suffice and indeed will be the default. Finally, you will be prompted for a password for this account.

If at any point after installation you would like to create another account, use the **adduser** command.

6.3.4. Installing the Base System

Although this stage is the least problematic, it consumes a significant fraction of the install because it downloads, verifies and unpacks the whole base system. If you have a slow computer or network connection, this could take some time.

During installation of the base system, package unpacking and setup messages are redirected to `tty4`. You can access this terminal by pressing **Left Alt-F4**; get back to the main installer process with **Left Alt-F1**.

The unpack/setup messages generated during this phase are also saved in `/var/log/syslog`. You can check them there if the installation is performed over a serial console.

As part of the installation, a Linux kernel will be installed. At the default priority, the installer will choose one for you that best matches your hardware. In lower priority modes, you will be able to choose from a list of available kernels.

6.3.5. Installing Additional Software

After the base system is installed, you have a usable but limited system. Most users will want to install additional software on the system to tune it to their needs, and the installer allows you do so. This step can take even longer than installing the base system if you have a slow computer or network.

6.3.5.1. Configuring apt

One of the tools used to install packages on a Debian GNU/Linux system is a program called **apt-get**, from the `apt` package[6]. Other front-ends for package management, like **aptitude** and **synaptic**, are also in use. These front-ends are recommended for new users, since they integrate some additional features (package searching and status checks) in a nice user interface. In fact, **aptitude** is now the recommended utility for package management.

apt must be configured so that it knows where to retrieve packages from. The installer largely takes care of this automatically based on what it knows about your installation medium. The results of this configuration are written to the file `/etc/apt/sources.list`, and you can examine and edit it to your liking after the install is complete.

6.3.5.2. Selecting and Installing Software

During the installation process, you are given the opportunity to select additional software to install. Rather than picking individual software packages from the 17600 available packages, this stage of the

6. Note that the program which actually installs the packages is called **dpkg**. However, this program is more of a low-level tool. **apt-get** is a higher-level tool, which will invoke **dpkg** as appropriate. It knows how to retrieve packages from your CD, the network, or wherever. It is also able to automatically install other packages which are required to make the package you're trying to install work correctly.

installation process focuses on selecting and installing predefined collections of software to quickly set up your computer to perform various tasks.

So, you have the ability to choose *tasks* first, and then add on more individual packages later. These tasks loosely represent a number of different jobs or things you want to do with your computer, such as "Desktop environment", "Web server", or "Print server"[7]. Section D.2 lists the space requirements for the available tasks.

Some tasks may be pre-selected based on the characteristics of the computer you are installing. If you disagree with these selections you can un-select the tasks. You can even opt to install no tasks at all at this point.

> **Note:** The "Desktop environment" task will install the GNOME desktop environment. The options offered by the installer currently do not allow to select a different desktop environment such as for example KDE.
>
> It is possible to get the installer to install KDE by using preseeding (see Section B.4.11) or by adding `tasks="standard, kde-desktop"` at the boot prompt when starting the installer. However, this will only work if the packages needed for KDE are actually available. If you are installing using a full CD image, they will need to be downloaded from a mirror as KDE packages are not included on the first full CD; installing KDE this way should work fine if you are using a DVD image or any other installation method.
>
> The various server tasks will install software roughly as follows. DNS server: `bind9`; File server: `samba`, `nfs`; Mail server: `exim4`, `spamassassin`, `uw-imap`; Print server: `cups`; SQL database: `postgresql`; Web server: `apache`.

Once you've selected your tasks, select OK. At this point, **aptitude** will install the packages that are part of the tasks you've selected.

> **Note:** In the standard user interface of the installer, you can use the space bar to toggle selection of a task.

You should be aware that especially the Desktop task is very large. Especially when installing from a normal CD-ROM in combination with a mirror for packages not on the CD-ROM, the installer may want to retrieve a lot of packages over the network. If you have a relatively slow Internet connection, this can take a long time. There is no option to cancel the installation of packages once it has started.

Even when packages are included on the CD-ROM, the installer may still retrieve them from the mirror if the version available on the mirror is more recent than the one included on the CD-ROM. If you are installing the stable distribution, this can happen after a point release (an update of the original stable release); if you are installing the testing distribution this will happen if you are using an older image.

Each package you selected with **tasksel** is downloaded, unpacked and then installed in turn by the **apt-get** and **dpkg** programs. If a particular program needs more information from the user, it will prompt you during this process.

7. You should know that to present this list, the installer is merely invoking the **tasksel** program. It can be run at any time after installation to install more packages (or remove them), or you can use a more fine-grained tool such as **aptitude**. If you are looking for a specific single package, after installation is complete, simply run **`aptitude install`** ***`package`***, where *`package`* is the name of the package you are looking for.

6.3.6. Making Your System Bootable

If you are installing a diskless workstation, obviously, booting off the local disk isn't a meaningful option, and this step will be skipped.

6.3.6.1. Detecting other operating systems

Before a boot loader is installed, the installer will attempt to probe for other operating systems which are installed on the machine. If it finds a supported operating system, you will be informed of this during the boot loader installation step, and the computer will be configured to boot this other operating system in addition to Debian.

Note that multiple operating systems booting on a single machine is still something of a black art. The automatic support for detecting and setting up boot loaders to boot other operating systems varies by architecture and even by subarchitecture. If it does not work you should consult your boot manager's documentation for more information.

6.3.6.2. palo-installer

The bootloader on PA-RISC is "palo". **PALO** is similar in configuration and usage to **LILO**, with a few exceptions. First of all, **PALO** allows you to boot any kernel image on your boot partition. This is because **PALO** can actually read Linux partitions.

6.3.6.3. Continue Without Boot Loader

This option can be used to complete the installation even when no boot loader is to be installed, either because the arch/subarch doesn't provide one, or because none is desired (e.g. you will use existing boot loader).

If you plan to manually configure your bootloader, you should check the name of the installed kernel in `/target/boot`. You should also check that directory for the presence of an *initrd*; if one is present, you will probably have to instruct your bootloader to use it. Other information you will need are the disk and partition you selected for your `/` filesystem and, if you chose to install `/boot` on a separate partition, also your `/boot` filesystem.

6.3.7. Finishing the Installation

These are the last bits to do before rebooting to your new system. It mostly consists of tidying up after the `debian-installer`.

6.3.7.1. Finish the Installation and Reboot

This is the last step in the initial Debian installation process. You will be prompted to remove the boot media (CD, floppy, etc) that you used to boot the installer. The installer will do any last minute tasks, and then reboot into your new Debian system.

6.3.8. Miscellaneous

The components listed in this section are usually not involved in the installation process, but are waiting in the background to help the user in case something goes wrong.

6.3.8.1. Saving the installation logs

If the installation is successful, the logfiles created during the installation process will be automatically saved to `/var/log/installer/` on your new Debian system.

Choosing Save debug logs from the main menu allows you to save the log files to a floppy disk, network, hard disk, or other media. This can be useful if you encounter fatal problems during the installation and wish to study the logs on another system or attach them to an installation report.

6.3.8.2. Using the Shell and Viewing the Logs

There are several methods you can use to get a shell while running an installation. On most systems, and if you are not installing over serial console, the easiest method is to switch to the second *virtual console* by pressing **Left Alt-F2**[8] (on a Mac keyboard, **Option-F2**). Use **Left Alt-F1** to switch back to installer itself.

If you cannot switch consoles, there is also an Execute a Shell item on the main menu that can be used to start a shell. To get back to the installer itself, type **`exit`** to close the shell.

At this point you are booted from the RAM disk, and there is a limited set of Unix utilities available for your use. You can see what programs are available with the command **ls /bin /sbin /usr/bin /usr/sbin** and by typing **help**. The shell is a Bourne shell clone called **ash** and has some nice features like autocompletion and history.

To edit and view files, use the text editor **nano**. Log files for the installation system can be found in the `/var/log` directory.

> **Note:** Although you can do basically anything in a shell that the available commands allow you to do, the option to use a shell is really only there in case something goes wrong and for debugging.
>
> Doing things manually from the shell may interfere with the installation process and result in errors or an incomplete installation. In particular, you should always use let the installer activate your swap partition and not do this yourself from a shell.

6.3.8.3. Installation Over the Network

One of the more interesting components is *network-console*. It allows you to do a large part of the installation over the network via SSH. The use of the network implies you will have to perform the first steps of the installation from the console, at least to the point of setting up the networking. (Although you can automate that part with Section 4.4.)

This component is not loaded into the main installation menu by default, so you have to explicitly ask for it. If you are installing from CD, you need to boot with medium priority or otherwise invoke the main installation menu and choose Load installer components from CD and from the list of additional components select network-console: Continue installation remotely using SSH. Successful load is indicated by a new menu entry called Continue installation remotely using SSH.

8. That is: press the **Alt** key on the left-hand side of the **space bar** and the **F2** function key at the same time.

After selecting this new entry, you will be asked for a new password to be used for connecting to the installation system and for its confirmation. That's all. Now you should see a screen which instructs you to login remotely as the user *installer* with the password you just provided. Another important detail to notice on this screen is the fingerprint of this system. You need to transfer the fingerprint securely to the "person who will continue the installation remotely".

Should you decide to continue with the installation locally, you can always press **Enter**, which will bring you back to the main menu, where you can select another component.

Now let's switch to the other side of the wire. As a prerequisite, you need to configure your terminal for UTF-8 encoding, because that is what the installation system uses. If you do not, remote installation will be still possible, but you may encounter strange display artefacts like destroyed dialog borders or unreadable non-ascii characters. Establishing a connection with the installation system is as simple as typing:

```
$ ssh -l installer install_host
```

Where *`install_host`* is either the name or IP address of the computer being installed. Before the actual login the fingerprint of the remote system will be displayed and you will have to confirm that it is correct.

> **Note:** If you install several computers in turn and they happen to have the same IP address or hostname, **ssh** will refuse to connect to such host. The reason is that it will have different fingerprint, which is usually a sign of a spoofing attack. If you are sure this is not the case, you will need to delete the relevant line from `~/.ssh/known_hosts` and try again.

After the login you will be presented with an initial screen where you have two possibilities called Start menu and Start shell. The former brings you to the main installer menu, where you can continue with the installation as usual. The latter starts a shell from which you can examine and possibly fix the remote system. You should only start one SSH session for the installation menu, but may start multiple sessions for shells.

Warning

After you have started the installation remotely over SSH, you should not go back to the installation session running on the local console. Doing so may corrupt the database that holds the configuration of the new system. This in turn may result in a failed installation or problems with the installed system.

Also, if you are running the SSH session from an X terminal, you should not resize the window as that will result in the connection being terminated.

Chapter 7. Booting Into Your New Debian System

7.1. The Moment of Truth

Your system's first boot on its own power is what electrical engineers call the "smoke test".

If the system fails to start up correctly, don't panic. If the installation was successful, chances are good that there is only a relatively minor problem that is preventing the system from booting Debian. In most cases such problems can be fixed without having to repeat the installation. One available option to fix boot problems is to use the installer's built-in rescue mode (see Section 8.7).

If you are new to Debian and Linux, you may need some help from more experienced users. For less common architectures like PA-RISC, your best option is to ask on the debian-hppa mailing list (http://www.debian.org/MailingLists/subscribe). You can also file an installation report as described in Section 5.2.5. Please make sure that you describe your problem clearly and include any messages that are displayed and may help others to diagnose the issue.

7.2. Mounting encrypted volumes

If you created encrypted volumes during the installation and assigned them mount points, you will be asked to enter the passphrase for each of these volumes during the boot. The actual procedure differs slightly between dm-crypt and loop-AES.

7.2.1. dm-crypt

For partitions encrypted using dm-crypt you will be shown the following prompt during the boot:

```
Starting early crypto disks... part_crypt(starting)
Enter LUKS passphrase:
```

In the first line of the prompt, *part* is the name of the underlying partition, e.g. sda2 or md0. You are now probably wondering *for which volume* you are actually entering the passphrase. Does it relate to your `/home`? Or to `/var`? Of course, if you have just one encrypted volume, this is easy and you can just enter the passphrase you used when setting up this volume. If you set up more than one encrypted volume during the installation, the notes you wrote down as the last step in Section 6.3.2.4 come in handy. If you did not make a note of the mapping between *part*`_crypt` and the mount points before, you can still find it in `/etc/crypttab` and `/etc/fstab` of your new system.

The prompt may look somewhat different when an encrypted root file system is mounted. This depends on which initramfs generator was used to generate the initrd used to boot the system. The example below is for an initrd generated using `initramfs-tools`:

```
Begin: Mounting root file system... ...
Begin: Running /scripts/local-top ...
Enter LUKS passphrase:
```

No characters (even asterisks) will be shown while entering the passphrase. If you enter the wrong passphrase, you have two more tries to correct it. After the third try the boot process will skip this volume and continue to mount the next filesystem. Please see Section 7.2.3 for further information.

After entering all passphrases the boot should continue as usual.

7.2.2. loop-AES

For partitions encrypted using loop-AES you will be shown the following prompt during the boot:

```
Checking loop-encrypted file systems.
Setting up /dev/loopX (/mountpoint)
Password:
```

No characters (even asterisks) will be shown while entering the passphrase. If you enter the wrong passphrase, you have two more tries to correct it. After the third try the boot process will skip this volume and continue to mount the next filesystem. Please see Section 7.2.3 for further information.

After entering all passphrases the boot should continue as usual.

7.2.3. Troubleshooting

If some of the encrypted volumes could not be mounted because a wrong passphrase was entered, you will have to mount them manually after the boot. There are several cases.

- The first case concerns the root partition. When it is not mounted correctly, the boot process will halt and you will have to reboot the computer to try again.
- The easiest case is for encrypted volumes holding data like `/home` or `/srv`. You can simply mount them manually after the boot. For loop-AES this is one-step operation:

  ```
  # mount /mount_point
  Password:
  ```

  where *`/mount_point`* should be replaced by the particular directory (e.g. `/home`). The only difference from an ordinary mount is that you will be asked to enter the passphrase for this volume.

  For dm-crypt this is a bit trickier. First you need to register the volumes with device mapper by running:

  ```
  # /etc/init.d/cryptdisks start
  ```

  This will scan all volumes mentioned in `/etc/crypttab` and will create appropriate devices under the `/dev` directory after entering the correct passphrases. (Already registered volumes will be skipped, so you can repeat this command several times without worrying.) After successful registration you can simply mount the volumes the usual way:

  ```
  # mount /mount_point
  ```

- If any volume holding noncritical system files could not be mounted (`/usr` or `/var`), the system should still boot and you should be able to mount the volumes manually like in the previous case. However, you will also need to (re)start any services usually running in your default runlevel because it is very likely that they were not started. The easiest way to achieve this is by switching to the first runlevel and back by entering

  ```
  # init 1
  ```

  at the shell prompt and pressing **Control-D** when asked for the root password.

7.3. Log In

Once your system boots, you'll be presented with the login prompt. Log in using the personal login and password you selected during the installation process. Your system is now ready for use.

If you are a new user, you may want to explore the documentation which is already installed on your system as you start to use it. There are currently several documentation systems, work is proceeding on integrating the different types of documentation. Here are a few starting points.

Documentation accompanying programs you have installed can be found in `/usr/share/doc/`, under a subdirectory named after the program (or, more precise, the Debian package that contains the program). However, more extensive documentation is often packaged separately in special documentation packages that are mostly not installed by default. For example, documentation about the package management tool **apt** can be found in the packages `apt-doc` or `apt-howto`.

In addition, there are some special folders within the `/usr/share/doc/` hierarchy. Linux HOWTOs are installed in *.gz* (compressed) format, in `/usr/share/doc/HOWTO/en-txt/`. After installing `dhelp`, you will find a browse-able index of documentation in `/usr/share/doc/HTML/index.html`.

One easy way to view these documents using a text based browser is to enter the following commands:

```
$ cd /usr/share/doc/
$ w3c .
```

The dot after the **w3c** command tells it to show the contents of the current directory.

If you have a graphical desktop environment installed, you can also use its web browser. Start the web browser from the application menu and enter **`/usr/share/doc/`** in the address bar.

You can also type **`info`** ***`command`*** or **`man`** ***`command`*** to see documentation on most commands available at the command prompt. Typing **`help`** will display help on shell commands. And typing a command followed by **`--help`** will usually display a short summary of the command's usage. If a command's results scroll past the top of the screen, type **`| more`** after the command to cause the results to pause before scrolling past the top of the screen. To see a list of all commands available which begin with a certain letter, type the letter and then two tabs.

Chapter 8. Next Steps and Where to Go From Here

8.1. Shutting down the system

To shut down a running Linux system, you must not reboot with the reset switch on the front or back of your computer, or just turn off the computer. Linux should be shut down in a controlled manner, otherwise files might get lost and/or disk damage might occur. If you run a desktop environment, there is usually an option to "log out" available from the application menu that allows you to shutdown (or reboot) the system.

Alternatively you can press the key combination **Ctrl-Alt-Del** . A last option is to log in as root and type one of the commands **poweroff**, **halt** or **shutdown -h now** if either of the key combinations do not work or you prefer to type commands; use **reboot** to reboot the system.

8.2. If You Are New to Unix

If you are new to Unix, you probably should go out and buy some books and do some reading. A lot of valuable information can also be found in the Debian Reference (http://www.debian.org/doc/user-manuals#quick-reference). This list of Unix FAQs (http://www.faqs.org/faqs/unix-faq/) contains a number of UseNet documents which provide a nice historical reference.

Linux is an implementation of Unix. The Linux Documentation Project (LDP) (http://www.tldp.org/) collects a number of HOWTOs and online books relating to Linux. Most of these documents can be installed locally; just install the `doc-linux-html` package (HTML versions) or the `doc-linux-text` package (ASCII versions), then look in `/usr/share/doc/HOWTO`. International versions of the LDP HOWTOs are also available as Debian packages.

8.3. Orienting Yourself to Debian

Debian is a little different from other distributions. Even if you're familiar with Linux in other distributions, there are things you should know about Debian to help you to keep your system in a good, clean state. This chapter contains material to help you get oriented; it is not intended to be a tutorial for how to use Debian, but just a very brief glimpse of the system for the very rushed.

8.3.1. Debian Packaging System

The most important concept to grasp is the Debian packaging system. In essence, large parts of your system should be considered under the control of the packaging system. These include:

- `/usr` (excluding `/usr/local`)
- `/var` (you could make `/var/local` and be safe in there)
- `/bin`
- `/sbin`

- `/lib`

For instance, if you replace `/usr/bin/perl`, that will work, but then if you upgrade your `perl` package, the file you put there will be replaced. Experts can get around this by putting packages on "hold" in **aptitude**.

One of the best installation methods is apt. You can use the command line version **apt-get** or full-screen text version aptitude. Note apt will also let you merge main, contrib, and non-free so you can have export-restricted packages as well as standard versions.

8.3.2. Application Version Management

Alternative versions of applications are managed by update-alternatives. If you are maintaining multiple versions of your applications, read the update-alternatives man page.

8.3.3. Cron Job Management

Any jobs under the purview of the system administrator should be in `/etc`, since they are configuration files. If you have a root cron job for daily, weekly, or monthly runs, put them in `/etc/cron.{daily,weekly,monthly}`. These are invoked from `/etc/crontab`, and will run in alphabetic order, which serializes them.

On the other hand, if you have a cron job that (a) needs to run as a special user, or (b) needs to run at a special time or frequency, you can use either `/etc/crontab`, or, better yet, `/etc/cron.d/whatever`. These particular files also have an extra field that allows you to stipulate the user account under which the cron job runs.

In either case, you just edit the files and cron will notice them automatically. There is no need to run a special command. For more information see cron(8), crontab(5), and `/usr/share/doc/cron/README.Debian`.

8.4. Further Reading and Information

If you need information about a particular program, you should first try **`man`** ***`program`***, or **`info`** ***`program`***.

There is lots of useful documentation in `/usr/share/doc` as well. In particular, `/usr/share/doc/HOWTO` and `/usr/share/doc/FAQ` contain lots of interesting information. To submit bugs, look at `/usr/share/doc/debian/bug*`. To read about Debian-specific issues for particular programs, look at `/usr/share/doc/(package name)/README.Debian`.

The Debian web site (http://www.debian.org/) contains a large quantity of documentation about Debian. In particular, see the Debian GNU/Linux FAQ (http://www.debian.org/doc/FAQ/) and the Debian Reference (http://www.debian.org/doc/user-manuals#quick-reference). An index of more Debian documentation is available from the Debian Documentation Project (http://www.debian.org/doc/ddp). The Debian community is self-supporting; to subscribe to one or more of the Debian mailing lists, see the Mail List Subscription (http://www.debian.org/MailingLists/subscribe) page. Last, but not least, the Debian Mailing List Archives (http://lists.debian.org/) contain a wealth of information on Debian.

A general source of information on GNU/Linux is the Linux Documentation Project (http://www.tldp.org/). There you will find the HOWTOs and pointers to other very valuable information on parts of a GNU/Linux system.

8.5. Setting Up Your System To Use E-Mail

Today, email is an important part of many people's life. As there are many options as to how to set it up, and as having it set up correctly is important for some Debian utilities, we will try to cover the basics in this section.

There are three main functions that make up an e-mail system. First there is the *Mail User Agent* (MUA) which is the program a user actually uses to compose and read mails. Then there is the *Mail Transfer Agent* (MTA) that takes care of transferring messages from one computer to another. And last there is the *Mail Delivery Agent* (MDA) that takes care of delivering incoming mail to the user's inbox.

These three functions can be performed by separate programs, but they can also be combined in one or two programs. It is also possible to have different programs handle these functions for different types of mail.

On Linux and Unix systems **mutt** is historically a very popular MUA. Like most traditional Linux programs it is text based. It is often used in combination with **exim** or **sendmail** as MTA and **procmail** as MDA.

With the increasing popularity of graphical desktop systems, the use of graphical e-mail programs like GNOME's **evolution**, KDE's **kmail** or Mozilla's **thunderbird** (in Debian available as **icedove**[1]) is becoming more popular. These programs combine the function of a MUA, MTA and MDA, but can — and often are — also be used in combination with the traditional Linux tools.

8.5.1. Default E-Mail Configuration

Even if you are planning to use a graphical mail program, it is important that a traditional MTA/MDA is also installed and correctly set up on your Linux system. Reason is that various utilities running on the system[2] can send important notices by e-mail to inform the system administrator of (potential) problems or changes.

For this reason the packages `exim4` and `mutt` will be installed by default (provided you did not unselect the "standard" task during the installation). `exim4` is a combination MTA/MDA that is relatively small but very flexible. By default it will be configured to only handle e-mail local to the system itself and e-mails addressed to the system administrator (root account) will be delivered to the regular user account created during the installation[3].

When system e-mails are delivered they are added to a file in `/var/mail/`*`account_name`*. The e-mails can be read using **mutt**.

8.5.2. Sending E-Mails Outside The System

As mentioned earlier, the installed Debian system is only set up to handle e-mail local to the system, not for sending mail to others nor for receiving mail from others.

If you would like `exim4` to handle external e-mail, please refer to the next subsection for the basic available configuration options. Make sure to test that mail can be sent and received correctly.

If you intend to use a graphical mail program and use a mail server of your Internet Service Provider (ISP) or your company, there is not really any need to configure `exim4` for handling external e-mail.

1. The reason that **thunderbird** has been renamed to **icedove** in Debian has to do with licencing issues. Details are outside the scope of this manual.
2. Examples are: **cron**, **quota**, **logcheck**, **aide**, ...
3. The forwarding of mail for root to the regular user account is configured in `/etc/aliases`. If no regular user account was created, the mail will of course be delivered to the root account itself.

Just configure your favorite graphical mail program to use the correct servers to send and receive e-mail (how is outside the scope of this manual).

However, in that case you may need to configure individual utilities to correctly send e-mails. One such utility is **reportbug**, a program that facilitates submitting bug reports against Debian packages. By default it expects to be able to use `exim4` to submit bug reports.

To correctly set up **reportbug** to use an external mail server, please run the command **reportbug --configure** and answer "no" to the question if an MTA is available. You will then be asked for the SMTP server to be used for submitting bug reports.

8.5.3. Configuring the Exim4 Mail Transport Agent

If you would like your system to also handle external e-mail, you will need to reconfigure the `exim4` package[4]:

```
# dpkg-reconfigure exim4-config
```

After entering that command (as root), you will be asked if you want split the configuration into small files. If you are unsure, select the default option.

Next you will be presented with several common mail scenarios. Choose the one that most closely resembles your needs.

internet site

Your system is connected to a network and your mail is sent and received directly using SMTP. On the following screens you will be asked a few basic questions, like your machine's mail name, or a list of domains for which you accept or relay mail.

mail sent by smarthost

In this scenario your outgoing mail is forwarded to another machine, called a "smarthost", which takes care of sending the message on to its destination. The smarthost also usually stores incoming mail addressed to your computer, so you don't need to be permanently online. That also means you have to download your mail from the smarthost via programs like fetchmail.

In a lot of cases the smarthost will be your ISP's mail server, which makes this option very suitable for dial-up users. It can also be a company mail server, or even another system on your own network.

mail sent by smarthost; no local mail

This option is basically the same as the previous one except that the system will not be set up to handle mail for a local e-mail domain. Mail on the system itself (e.g. for the system administrator) will still be handled.

local delivery only

This is the option your system is configured for by default.

4. You can of course also remove `exim4` and replace it with an alternative MTA/MDA.

no configuration at this time

Choose this if you are absolutely convinced you know what you are doing. This will leave you with an unconfigured mail system — until you configure it, you won't be able to send or receive any mail and you may miss some important messages from your system utilities.

If none of these scenarios suits your needs, or if you need a finer grained setup, you will need to edit configuration files under the `/etc/exim4` directory after the installation is complete. More information about `exim4` may be found under `/usr/share/doc/exim4`; the file `README.Debian.gz` has further details about configuring `exim4` and explains where to find additional documentation.

Note that sending mail directly to the Internet when you don't have an official domain name, can result in your mail being rejected because of anti-spam measures on receiving servers. Using your ISP's mail server is preferred. If you still do want to send out mail directly, you may want to use a different e-mail address than is generated by default. If you use `exim4` as your MTA, this is possible by adding an entry in `/etc/email-addresses`.

8.6. Compiling a New Kernel

Why would someone want to compile a new kernel? It is often not necessary since the default kernel shipped with Debian handles most configurations. Also, Debian often offers several alternative kernels. So you may want to check first if there is an alternative kernel image package that better corresponds to your hardware. However, it can be useful to compile a new kernel in order to:

- handle special hardware needs, or hardware conflicts with the pre-supplied kernels
- use options of the kernel which are not supported in the pre-supplied kernels (such as high memory support)
- optimize the kernel by removing useless drivers to speed up boot time
- create a monolithic instead of a modularized kernel
- run an updated or development kernel
- learn more about linux kernels

8.6.1. Kernel Image Management

Don't be afraid to try compiling the kernel. It's fun and profitable.

To compile a kernel the Debian way, you need some packages: `fakeroot`, `kernel-package`, `linux-source-2.6` and a few others which are probably already installed (see `/usr/share/doc/kernel-package/README.gz` for the complete list).

This method will make a .deb of your kernel source, and, if you have non-standard modules, make a synchronized dependent .deb of those too. It's a better way to manage kernel images; `/boot` will hold the kernel, the System.map, and a log of the active config file for the build.

Note that you don't *have* to compile your kernel the "Debian way"; but we find that using the packaging system to manage your kernel is actually safer and easier. In fact, you can get your kernel sources right from Linus instead of `linux-source-2.6`, yet still use the `kernel-package` compilation method.

Note that you'll find complete documentation on using `kernel-package` under `/usr/share/doc/kernel-package`. This section just contains a brief tutorial.

Hereafter, we'll assume you have free rein over your machine and will extract your kernel source to somewhere in your home directory[5]. We'll also assume that your kernel version is 2.6.18. Make sure you are in the directory to where you want to unpack the kernel sources, extract them using **`tar xjf /usr/src/linux-source-2.6.18.tar.bz2`** and change to the directory `linux-source-2.6.18` that will have been created.

Now, you can configure your kernel. Run **`make xconfig`** if X11 is installed, configured and being run; run **`make menuconfig`** otherwise (you'll need `libncurses5-dev` installed). Take the time to read the online help and choose carefully. When in doubt, it is typically better to include the device driver (the software which manages hardware peripherals, such as Ethernet cards, SCSI controllers, and so on) you are unsure about. Be careful: other options, not related to a specific hardware, should be left at the default value if you do not understand them. Do not forget to select "Kernel module loader" in "Loadable module support" (it is not selected by default). If not included, your Debian installation will experience problems.

Clean the source tree and reset the `kernel-package` parameters. To do that, do **`make-kpkg clean`**.

Now, compile the kernel: **`fakeroot make-kpkg --initrd --revision=custom.1.0 kernel_image`**. The version number of "1.0" can be changed at will; this is just a version number that you will use to track your kernel builds. Likewise, you can put any word you like in place of "custom" (e.g., a host name). Kernel compilation may take quite a while, depending on the power of your machine.

Once the compilation is complete, you can install your custom kernel like any package. As root, do **`dpkg -i ../linux-image-2.6.18-`*`subarchitecture`*`_custom.1.0_hppa.deb`**. The *`subarchitecture`* part is an optional sub-architecture, depending on what kernel options you set. **`dpkg -i`** will install the kernel, along with some other nice supporting files. For instance, the `System.map` will be properly installed (helpful for debugging kernel problems), and `/boot/config-2.6.18` will be installed, containing your current configuration set. Your new kernel package is also clever enough to automatically update your boot loader to use the new kernel. If you have created a modules package, you'll need to install that package as well.

It is time to reboot the system: read carefully any warning that the above step may have produced, then **`shutdown -r now`**.

For more information on Debian kernels and kernel compilation, see the Debian Linux Kernel Handbook (http://kernel-handbook.alioth.debian.org/). For more information on `kernel-package`, read the fine documentation in `/usr/share/doc/kernel-package`.

8.7. Recovering a Broken System

Sometimes, things go wrong, and the system you've carefully installed is no longer bootable. Perhaps the boot loader configuration broke while trying out a change, or perhaps a new kernel you installed won't boot, or perhaps cosmic rays hit your disk and flipped a bit in `/sbin/init`. Regardless of the cause, you'll need to have a system to work from while you fix it, and rescue mode can be useful for this.

To access rescue mode, type **`rescue`** at the `boot:` prompt, or boot with the **`rescue/enable=true`** boot parameter. You'll be shown the first few screens of the installer, with a note in the corner of

5. There are other locations where you can extract kernel sources and build your custom kernel, but this is easiest as it does not require special permissions.

the display to indicate that this is rescue mode, not a full installation. Don't worry, your system is not about to be overwritten! Rescue mode simply takes advantage of the hardware detection facilities available in the installer to ensure that your disks, network devices, and so on are available to you while repairing your system.

Instead of the partitioning tool, you should now be presented with a list of the partitions on your system, and asked to select one of them. Normally, you should select the partition containing the root file system that you need to repair. You may select partitions on RAID and LVM devices as well as those created directly on disks.

If possible, the installer will now present you with a shell prompt in the file system you selected, which you can use to perform any necessary repairs.

If the installer cannot run a usable shell in the root file system you selected, perhaps because the file system is corrupt, then it will issue a warning and offer to give you a shell in the installer environment instead. You may not have as many tools available in this environment, but they will often be enough to repair your system anyway. The root file system you selected will be mounted on the `/target` directory.

In either case, after you exit the shell, the system will reboot.

Finally, note that repairing broken systems can be difficult, and this manual does not attempt to go into all the things that might have gone wrong or how to fix them. If you have problems, consult an expert.

Appendix A. Installation Howto

This document describes how to install Debian GNU/Linux etch for the PA-RISC ("hppa") with the new `debian-installer`. It is a quick walkthrough of the installation process which should contain all the information you will need for most installs. When more information can be useful, we will link to more detailed explanations in other parts of this document.

A.1. Preliminaries

If you encounter bugs during your install, please refer to Section 5.2.5 for instructions on how to report them. If you have questions which cannot be answered by this document, please direct them to the debian-boot mailing list (debian-boot@lists.debian.org) or ask on IRC (#debian-boot on the OFTC network).

A.2. Booting the installer

The debian-cd team provides builds of CD images using `debian-installer` on the Debian CD page (http://www.debian.org/CD/). For more information on where to get CDs, see Section 4.1.

Some installation methods require other images than CD images. Section 4.2.1 explains how to find images on Debian mirrors.

The subsections below will give the details about which images you should get for each possible means of installation.

A.2.1. CDROM

There are two different netinst CD images which can be used to install etch with the `debian-installer`. These images are intended to boot from CD and install additional packages over a network, hence the name 'netinst'. The difference between the two images is that on the full netinst image the base packages are included, whereas you have to download these from the web if you are using the business card image. If you'd rather, you can get a full size CD image which will not need the network to install. You only need the first CD of the set.

Download whichever type you prefer and burn it to a CD.

A.2.2. Booting from network

It's also possible to boot `debian-installer` completely from the net. The various methods to netboot depend on your architecture and netboot setup. The files in `netboot/` can be used to netboot `debian-installer`.

A.2.3. Booting from hard disk

It's possible to boot the installer using no removable media, but just an existing hard disk, which can have a different OS on it. Download `hd-media/initrd.gz`, `hd-media/vmlinuz`, and a Debian

CD image to the top-level directory of the hard disk. Make sure that the CD image has a filename ending in `.iso`. Now it's just a matter of booting linux with the initrd.

A.3. Installation

Once the installer starts, you will be greeted with an initial screen. Press **Enter** to boot, or read the instructions for other boot methods and parameters (see Section 5.1).

After a while you will be asked to select your language. Use the arrow keys to pick a language and press **Enter** to continue. Next you'll be asked to select your country, with the choices including countries where your language is spoken. If it's not on the short list, a list of all the countries in the world is available.

You may be asked to confirm your keyboard layout. Choose the default unless you know better.

Now sit back while debian-installer detects some of your hardware, and loads the rest of itself from CD, floppy, USB, etc.

Next the installer will try to detect your network hardware and set up networking by DHCP. If you are not on a network or do not have DHCP, you will be given the opportunity to configure the network manually.

Now it is time to partition your disks. First you will be given the opportunity to automatically partition either an entire drive, or available free space on a drive (guided partitioning). This is recommended for new users or anyone in a hurry. If you do not want to autopartition, choose Manual from the menu.

On the next screen you will see your partition table, how the partitions will be formatted, and where they will be mounted. Select a partition to modify or delete it. If you did automatic partitioning, you should just be able to choose Finish partitioning and write changes to disk from the menu to use what it set up. Remember to assign at least one partition for swap space and to mount a partition on `/`. Appendix C has more information about partitioning.

Now `debian-installer` formats your partitions and starts to install the base system, which can take a while. That is followed by installing a kernel.

The next steps are setting up your time zone and clock. The installer will try to select the correct settings automatically and will only ask if it cannot. This is followed by setting up user accounts. By default you will need to provide a password for the "root" (administrator) account and information necessary to create one regular user account.

The base system that was installed earlier is a working, but very minimal installation. To make the system more functional the next step allows you to install additional packages by selecting tasks. Before packages can be installed `apt` needs to be configured as that defines from where the packages will be retrieved. The "Standard system" task will be selected by default and should normally be installed. Select the "Desktop environment" task if you would like to have a graphical desktop after the installation. See Section 6.3.5.2 for additional information about this step.

The last step is to install a boot loader. If the installer detects other operating systems on your computer, it will add them to the boot menu and let you know.

`debian-installer` will now tell you that the installation has finished. Remove the cdrom or other boot media and hit **Enter** to reboot your machine. It should boot up into the newly installed system and allow you to log in. This is explained in Chapter 7.

If you need more information on the install process, see Chapter 6.

A.4. Send us an installation report

If you successfully managed an installation with `debian-installer`, please take time to provide us with a report. The simplest way to do so is to install the reportbug package (**aptitude install reportbug**), configure `reportbug` as explained in Section 8.5.2, and run **reportbug installation-reports**.

If you did not complete the install, you probably found a bug in debian-installer. To improve the installer it is necessary that we know about them, so please take the time to report them. You can use an installation report to report problems; if the install completely fails, see Section 5.2.4.

A.5. And finally…

We hope that your Debian installation is pleasant and that you find Debian useful. You might want to read Chapter 8.

Appendix B. Automating the installation using preseeding

This appendix explains how to preseed answers to questions in `debian-installer` to automate your installation.

The configuration fragments used in this appendix are also available as an example preconfiguration file from http://www.debian.org/releases/etch/example-preseed.txt.

B.1. Introduction

Preseeding provides a way to set answers to questions asked during the installation process, without having to manually enter the answers while the installation is running. This makes it possible to fully automate most types of installation and even offers some features not available during normal installations.

B.1.1. Preseeding methods

There are three methods that can be used for preseeding: *initrd*, *file* and *network*. Initrd preseeding will work with any installation method and supports preseeding of more things, but it requires the most preparation. File and network preseeding each can be used with different installation methods.

The following table shows which preseeding methods can be used with which installation methods.

Installation method	initrd	file	network
CD/DVD	yes	yes	yes[a]
netboot	yes	no	yes
hd-media	yes	yes	yes[a]
Notes: a. but only if you have network access, and set `preseed/url` appropriately			

An important difference between the preseeding methods is the point at which the preconfiguration file is loaded and processed. For initrd preseeding this is right at the start of the installation, before the first question is even asked. For file preseeding this is after the CD or CD image has been loaded. For network preseeding it is only after the network has been configured.

Obviously, any questions that have been processed before the preconfiguration file is loaded cannot be preseeded (this will include questions that are only displayed at medium or low priority, like the first hardware detection run). Section B.2.2 offers a way to avoid these questions being asked.

In order to avoid the questions that would normally appear before the preseeding occurs, you can start the installer in “auto” mode. This delays questions that would normally be asked too early for preseeding (i.e. language, country and keyboard selection) until after the network comes up, thus allowing them to be preseeded. It also runs the installation at critical priority, which avoids many unimportant questions. See Section B.2.3 for details.

B.1.2. Limitations

Although most questions used by `debian-installer` can be preseeded using this method, there are some notable exceptions. You must (re)partition an entire disk or use available free space on a disk; it is not possible to use existing partitions.

B.2. Using preseeding

You will first need to create a preconfiguration file and place it in the location from where you want to use it. Creating the preconfiguration file is covered later in this appendix. Putting it in the correct location is fairly straightforward for network preseeding or if you want to read the file off a floppy or usb-stick. If you want to include the file on a CD or DVD, you will have to remaster the ISO image. How to get the preconfiguration file included in the initrd is outside the scope of this document; please consult the developers' documentation for `debian-installer`.

An example preconfiguration file that you can use as basis for your own preconfiguration file is available from http://www.debian.org/releases/etch/example-preseed.txt. This file is based on the configuration fragments included in this appendix.

B.2.1. Loading the preconfiguration file

If you are using initrd preseeding, you only have to make sure a file named `preseed.cfg` is included in the root directory of the initrd. The installer will automatically check if this file is present and load it.

For the other preseeding methods you need to tell the installer what file to use when you boot it. This is normally done by passing the kernel a boot parameter, either manually at boot time or by editing the bootloader configuration file (e.g. `syslinux.cfg`) and adding the parameter to the end of the append line(s) for the kernel.

If you do specify the preconfiguration file in the bootloader configuration, you might change the configuration so you don't need to hit enter to boot the installer. For syslinux this means setting the timeout to `1` in `syslinux.cfg`.

To make sure the installer gets the right preconfiguration file, you can optionally specify a checksum for the file. Currently this needs to be a md5sum, and if specified it must match the preconfiguration file or the installer will refuse to use it.

```
Boot parameters to specify:
- if you're netbooting:
  preseed/url=http://host/path/to/preseed.cfg
  preseed/url/checksum=5da499872becccfeda2c4872f9171c3d

- if you're booting a remastered CD:
  preseed/file=/cdrom/preseed.cfg
  preseed/file/checksum=5da499872becccfeda2c4872f9171c3d

- if you're installing from USB media (put the preconfiguration file in the
  toplevel directory of the USB stick):
  preseed/file=/hd-media/preseed.cfg
  preseed/file/checksum=5da499872becccfeda2c4872f9171c3d
```

Note that `preseed/url` can be shortened to just `url` and `preseed/file` to just `file` when they are passed as boot parameters.

B.2.2. Using boot parameters to preseed questions

If a preconfiguration file cannot be used to preseed some steps, the install can still be fully automated, since you can pass preseed values on the command line when booting the installer.

Boot parameters can also be used if you do not really want to use preseeding, but just want to provide an answer for a specific question. Some examples where this can be useful are documented elsewhere in this manual.

To set a value to be used inside `debian-installer`, just pass ***`path/to/variable=value`*** for any of the preseed variables listed in the examples in this appendix. If a value is to be used to configure packages for the target system, you will need to prepend the *owner*[1] of the variable as in ***`owner:path/to/variable=value`***. If you don't specify the owner, the value for the variable will not be copied to the debconf database in the target system and thus remain unused during the configuration of the relevant package.

Note that some variables that are frequently set at the boot prompt have a shorter alias. If an alias is available, it is used in the examples in this appendix instead of the full variable. In particular, the `preseed/url` variable has been aliased as `url` and has some extra magic to allow shortened urls to be used. Another example is the `tasks` alias, which translates to `tasksel:tasksel/first`.

A "--" in the boot options has special meaning. Kernel parameters that appear after the last "--" may be copied into the bootloader configuration for the installed system (if supported by the installer for the bootloader). The installer will automatically filter out any options (like preconfiguration options) that it recognizes.

> **Note:** Current linux kernels (2.6.9 and later) accept a maximum of 32 command line options and 32 environment options, including any options added by default for the installer. If these numbers are exceeded, the kernel will panic (crash). (For earlier kernels, these numbers were lower.)

For most installations some of the default options in your bootloader configuration file, like `vga=normal`, may be safely removed which may allow you to add more options for preseeding.

> **Note:** It may not always be possible to specify values with spaces for boot parameters, even if you delimit them with quotes.

B.2.3. Auto mode

There are several features of Debian Installer that combine to allow fairly simple command lines at the boot prompt to result in arbitrarily complex customized automatic installs. To illustrate this, here are some examples that can be used at the boot prompt:

```
auto url=autoserver
```

1. The owner of a debconf variable (or template) is normally the name of the package that contains the corresponding debconf template. For variables used in the installer itself the owner is "d-i". Templates and variables can have more than one owner which helps to determine whether they can be removed from the debconf database if the package is purged.

This relies on there being a DHCP server that will get the machine to the point where `autoserver` can be resolved by DNS, perhaps after adding the local domain if that was provided by DHCP. If this was done at a site where the domain is `example.com`, and they have a reasonably sane DHCP setup, it would result in the preseed file being retrieved from `http://autoserver.example.com/d-i/etch/./preseed.cfg`.

The last part of that url (`d-i/etch/./preseed.cfg`) is taken from `auto-install/defaultroot`. By default this includes the directory `etch` to allow future versions to specify their own codename and let people migrate forwards in a controlled manner. The `/./` bit is used to indicate a root, relative to which subsequent paths can be anchored (for use in preseed/include and preseed/run). This allows files to be specified either as full URLs, paths starting with / that are thus anchored, or even paths relative to the location where the last preseed file was found. This can be used to construct more portable scripts where an entire hierarchy of scripts can be moved to a new location without breaking it, for example copying the files onto a USB stick when they started out on a web server. In this example, if the preseed file sets `preseed/run` to `/scripts/late_command.sh` then the file will be fetched from `http://autoserver.example.com/d-i/etch/./scripts/late_command.sh`.

If there is no local DHCP or DNS infrastructure, or if you do not want to use the default path to `preseed.cfg`, you can still use an explicit url, and if you don't use the `/./` element it will be anchored to the start of the path (i.e. the third `/` in the URL). Here is an example that requires minimal support from the local network infrastructure:

```
auto url=http://192.168.1.2/path/to/mypreseed.file
```

The way this works is that:

- if the URL is missing a protocol, http is assumed,
- if the hostname section contains no periods, it has the domain derived from DHCP appended to it, and
- if there's no `/`'s after the hostname, then the default path is added.

In addition to specifying the url, you can also specify settings that do not directly affect the behavior of `debian-installer` itself, but can be passed through to scripts specified using `preseed/run` in the loaded preseed file. At present, the only example of this is `auto-install/classes`, which has an alias `classes`. This can be used thus:

```
auto url=example.com classes=class_A;class_B
```

The classes could for example denote the type of system to be installed, or the localization to be used.

It is of course possible to extend this concept, and if you do, it is reasonable to use the auto-install namespace for this. So one might have something like `auto-install/style` which is then used in your scripts. If you feel the need to do this, please mention it on the <`debian-boot@lists.debian.org`> mailing list so that we can avoid namespace conflicts, and perhaps add an alias for the parameter for you.

The `auto` boot label is not yet defined on all architectures. The same effect may be achieved by simply adding the two parameters `auto=true priority=critical` to the kernel command line. The `auto` parameter is an alias for `auto-install/enable` and controls the delay of the locale and keyboard questions until after there has been a chance to preseed them, while `priority` is an alias

for `debconf/priority` and setting it to `critical` stops any questions with a lower priority from being asked.

Additional options that may be of interest while attempting to automate an install while using DHCP are: `interface=auto netcfg/dhcp_timeout=60` which makes the machine choose the first viable NIC and be more patient about getting a reply to its DHCP query.

> **Tip:** An extensive example of how to use this framework, including example scripts and classes, can be found on the website of its developer (http://hands.com/d-i/). The examples available there also show many other nice effects that can be achieved by creative use of preconfiguration.

B.2.4. Aliases useful with preseeding

The following aliases can be useful when using (auto mode) preseeding.

auto	auto-install/enable
classes	auto-install/classes
fb	debian-installer/framebuffer
locale	debian-installer/locale
priority	debconf/priority
file	preseed/file
url	preseed/url
interface	netcfg/choose_interface
hostname	netcfg/get_hostname
domain	netcfg/get_domain
protocol	mirror/protocol
suite	mirror/suite

B.2.5. Using a DHCP server to specify preconfiguration files

It's also possible to use DHCP to specify a preconfiguration file to download from the network. DHCP allows specifying a filename. Normally this is a file to netboot, but if it appears to be an URL then installation media that support network preseeding will download the file from the URL and use it as a preconfiguration file. Here is an example of how to set it up in the dhcpd.conf for version 3 of the ISC DHCP server (the dhcp3-server Debian package).

```
if substring (option vendor-class-identifier, 0, 3) = "d-i" {
    filename "http://host/preseed.cfg";
}
```

Note that the above example limits this filename to DHCP clients that identify themselves as "d-i", so it will not affect regular DHCP clients, but only the installer. You can also put the text in a stanza for only one particular host to avoid preseeding all installs on your network.

A good way to use the DHCP preseeding is to only preseed values specific to your network, such as the Debian mirror to use. This way installs on your network will automatically get a good mirror

selected, but the rest of the installation can be performed interactively. Using DHCP preseeding to fully automate Debian installs should only be done with care.

B.3. Creating a preconfiguration file

The preconfiguration file is in the format used by the **debconf-set-selections** command. The general format of a line in a preconfiguration file is:

```
<owner> <question name> <question type> <value>
```

There are a few rules to keep in mind when writing a preconfiguration file.

- Put only a single space or tab between type and value: any additional whitespace will be interpreted as belonging to the value.
- A line can be split into multiple lines by appending a backslash ("\") as the line continuation character. A good place to split a line is after the question name; a bad place is between type and value.
- Most questions need to be preseeded using the values valid in English and not the translated values. However, there are some questions (for example in `partman`) where the translated values need to be used.
- Some questions take a code as value instead of the English text that is shown during installation.

The easiest way to create a preconfiguration file is to use the example file linked in Section B.4 as basis and work from there.

An alternative method is to do a manual installation and then, after rebooting, use the **debconf-get-selections** from the `debconf-utils` package to dump both the debconf database and the installer's cdebconf database to a single file:

```
$ debconf-get-selections --installer > file
$ debconf-get-selections >> file
```

However, a file generated in this manner will have some items that should not be preseeded, and the example file is a better starting place for most users.

> **Note:** This method relies on the fact that, at the end of the installation, the installer's cdebconf database is saved to the installed system in `/var/log/installer/cdebconf`. However, because the database may contain sensitive information, by default the files are only readable by root.
>
> The directory `/var/log/installer` and all files in it will be deleted from your system if you purge the package `installation-report`.

To check possible values for questions, you can use **nano** to examine the files in `/var/lib/cdebconf` while an installation is in progress. View `templates.dat` for the raw templates and `questions.dat` for the current values and for the values assigned to variables.

To check if the format of your preconfiguration file is valid before performing an install, you can use the command **debconf-set-selections -c** ***preseed.cfg***.

B.4. Contents of the preconfiguration file

The configuration fragments used in this appendix are also available as an example preconfiguration file from http://www.debian.org/releases/etch/example-preseed.txt.

Note that this example is based on an installation for the Intel x86 architecture. If you are installing a different architecture, some of the examples (like keyboard selection and bootloader installation) may not be relevant and will need to be replaced by debconf settings appropriate for your architecture.

B.4.1. Localization

Setting localization values will only work if you are using initrd preseeding. With all other methods the preconfiguration file will only be loaded after these questions have been asked.

The locale can be used to specify both language and country. To specify the locale as a boot parameter, use **locale=*en_US***.

```
# Locale sets language and country.
d-i debian-installer/locale string en_US
```

Keyboard configuration consists of selecting a keyboard architecture and a keymap. In most cases the correct keyboard architecture is selected by default, so there's normally no need to preseed it. The keymap must be valid for the selected keyboard architecture.

```
# Keyboard selection.
#d-i console-tools/archs select at
d-i console-keymaps-at/keymap select us
# Example for a different keyboard architecture
#d-i console-keymaps-usb/keymap select mac-usb-us
```

To skip keyboard configuration, preseed `console-tools/archs` with **`skip-config`**. This will result in the kernel keymap remaining active.

> **Note:** The changes in the input layer for 2.6 kernels have made the keyboard architecture virtually obsolete. For 2.6 kernels normally a "PC" (**`at`**) keymap should be selected.

B.4.2. Network configuration

Of course, preseeding the network configuration won't work if you're loading your preconfiguration file from the network. But it's great when you're booting from CD or USB stick. If you are loading preconfiguration files from the network, you can pass network config parameters by using kernel boot parameters.

If you need to pick a particular interface when netbooting before loading a preconfiguration file from the network, use a boot parameter such as **`interface=`*`eth1`***.

Although preseeding the network configuration is normally not possible when using network preseeding (using "preseed/url"), you can use the following hack to work around that, for example if you'd like to set a static address for the network interface. The hack is to force the network configuration to run again after the preconfiguration file has been loaded by creating a "preseed/run" script containing the following lines:

```
killall.sh dhclient
netcfg
```

```
# netcfg will choose an interface that has link if possible. This makes it
# skip displaying a list if there is more than one interface.
d-i netcfg/choose_interface select auto

# To pick a particular interface instead:
#d-i netcfg/choose_interface select eth1

# If you have a slow dhcp server and the installer times out waiting for
# it, this might be useful.
#d-i netcfg/dhcp_timeout string 60

# If you prefer to configure the network manually, uncomment this line and
# the static network configuration below.
#d-i netcfg/disable_dhcp boolean true

# If you want the preconfiguration file to work on systems both with and
# without a dhcp server, uncomment these lines and the static network
# configuration below.
#d-i netcfg/dhcp_failed note
#d-i netcfg/dhcp_options select Configure network manually

# Static network configuration.
#d-i netcfg/get_nameservers string 192.168.1.1
#d-i netcfg/get_ipaddress string 192.168.1.42
#d-i netcfg/get_netmask string 255.255.255.0
#d-i netcfg/get_gateway string 192.168.1.1
#d-i netcfg/confirm_static boolean true

# Any hostname and domain names assigned from dhcp take precedence over
# values set here. However, setting the values still prevents the questions
# from being shown, even if values come from dhcp.
d-i netcfg/get_hostname string unassigned-hostname
d-i netcfg/get_domain string unassigned-domain

# Disable that annoying WEP key dialog.
d-i netcfg/wireless_wep string
# The wacky dhcp hostname that some ISPs use as a password of sorts.
#d-i netcfg/dhcp_hostname string radish
```

B.4.3. Mirror settings

Depending on the installation method you use, a mirror may be used to download additional components of the installer, to install the base system, and to set up the `/etc/apt/sources.list` for the installed system.

The parameter `mirror/suite` determines the suite for the installed system.

The parameter `mirror/udeb/suite` determines the suite for additional components for the installer. It is only useful to set this if components are actually downloaded over the network and should match the suite that was used to build the initrd for the installation method used for the installation. By default the value for `mirror/udeb/suite` is the same as `mirror/suite`.

```
# If you select ftp, the mirror/country string does not need to be set.
#d-i mirror/protocol string ftp
d-i mirror/country string enter information manually
d-i mirror/http/hostname string http.us.debian.org
d-i mirror/http/directory string /debian
d-i mirror/http/proxy string

# Suite to install.
#d-i mirror/suite string testing
# Suite to use for loading installer components (optional).
#d-i mirror/udeb/suite string testing
```

B.4.4. Partitioning

Using preseeding to partition the harddisk is very much limited to what is supported by `partman-auto`. You can choose to partition either existing free space on a disk or a whole disk. The layout of the disk can be determined by using a predefined recipe, a custom recipe from a recipe file or a recipe included in the preconfiguration file. It is currently not possible to partition multiple disks using preseeding.

Warning

The identification of disks is dependent on the order in which their drivers are loaded. If there are multiple disks in the system, make very sure the correct one will be selected before using preseeding.

```
# If the system has free space you can choose to only partition that space.
# Note: this must be preseeded with a localized (translated) value.
#d-i partman-auto/init_automatically_partition \
#      select Guided - use the largest continuous free space

# Alternatively, you can specify a disk to partition. The device name
# can be given in either devfs or traditional non-devfs format.
# For example, to use the first disk:
d-i partman-auto/disk string /dev/discs/disc0/disc
# In addition, you'll need to specify the method to use.
# The presently available methods are: "regular", "lvm" and "crypto"
d-i partman-auto/method string lvm

# If one of the disks that are going to be automatically partitioned
# contains an old LVM configuration, the user will normally receive a
```

```
# warning. This can be preseeded away...
d-i partman-auto/purge_lvm_from_device boolean true
# And the same goes for the confirmation to write the lvm partitions.
d-i partman-lvm/confirm boolean true

# You can choose from any of the predefined partitioning recipes.
# Note: this must be preseeded with a localized (translated) value.
d-i partman-auto/choose_recipe \
       select All files in one partition (recommended for new users)
#d-i partman-auto/choose_recipe \
#       select Separate /home partition
#d-i partman-auto/choose_recipe \
#       select Separate /home, /usr, /var, and /tmp partitions

# Or provide a recipe of your own...
# The recipe format is documented in the file devel/partman-auto-recipe.txt.
# If you have a way to get a recipe file into the d-i environment, you can
# just point at it.
#d-i partman-auto/expert_recipe_file string /hd-media/recipe

# If not, you can put an entire recipe into the preconfiguration file in one
# (logical) line. This example creates a small /boot partition, suitable
# swap, and uses the rest of the space for the root partition:
#d-i partman-auto/expert_recipe string                         \
#      boot-root ::                                            \
#              40 50 100 ext3                                  \
#                      $primary{ } $bootable{ }                \
#                      method{ format } format{ }              \
#                      use_filesystem{ } filesystem{ ext3 }    \
#                      mountpoint{ /boot }                     \
#              .                                               \
#              500 10000 1000000000 ext3                       \
#                      method{ format } format{ }              \
#                      use_filesystem{ } filesystem{ ext3 }    \
#                      mountpoint{ / }                         \
#              .                                               \
#              64 512 300% linux-swap                          \
#                      method{ swap } format{ }                \
#              .

# This makes partman automatically partition without confirmation.
d-i partman/confirm_write_new_label boolean true
d-i partman/choose_partition \
       select Finish partitioning and write changes to disk
d-i partman/confirm boolean true
```

B.4.5. Partitioning using RAID

You can also use preseeding to set up partitions on software RAID arrays. Supported are RAID levels 0, 1 and 5, creating degraded arrays and specifying spare devices. If you are using RAID 1, you can preseed grub to install to all devices used in the array; see Section B.4.10.

Warning

This type of automated partitioning is easy to get wrong. It is also a very new component that may still have some bugs or missing error handling. The responsibility to get the various recipes right (so they make sense and don't conflict) lies with the user. Check `/var/log/syslog` if you run into problems.

Note that only RAID 0 and RAID 1 have been tested by the developers of the component. RAID 5 is untested. Advanced RAID setup with degraded arrays or spare devices has only been tested lightly.

```
# NOTE: this option is of beta release quality and should be used carefully

# The method should be set to "raid".
#d-i partman-auto/method string raid
# Specify the disks to be partitioned. They will all get the same layout,
# so this will only work if the disks are the same size.
#d-i partman-auto/disk string /dev/discs/disc0/disc /dev/discs/disc1/disc

# Next you need to specify the physical partitions that will be used.
#d-i partman-auto/expert_recipe string \
#      multiraid ::                                          \
#              1000 5000 4000 raid                           \
#                      $primary{ } method{ raid }            \
#              .                                             \
#              64 512 300% raid                              \
#                      method{ raid }                        \
#              .                                             \
#              500 10000 1000000000 raid                     \
#                      method{ raid }                        \
#              .

# Last you need to specify how the previously defined partitions will be
# used in the RAID setup. Remember to use the correct partition numbers
# for logical partitions.
# Parameters are:
# <raidtype> <devcount> <sparecount> <fstype> <mountpoint> \
#          <devices> <sparedevices>
# RAID levels 0, 1 and 5 are supported; devices are separated using "#"
#d-i partman-auto-raid/recipe string \
#    1 2 0 ext3 /                                            \
#          /dev/discs/disc0/part1#/dev/discs/disc1/part1     \
#    .                                                       \
#    1 2 0 swap -                                            \
#          /dev/discs/disc0/part5#/dev/discs/disc1/part5     \
#    .                                                       \
#    0 2 0 ext3 /home                                        \
#          /dev/discs/disc0/part6#/dev/discs/disc1/part6     \
#    .

# This makes partman automatically partition without confirmation.
d-i partman-md/confirm boolean true
d-i partman/confirm_write_new_label boolean true
d-i partman/choose_partition \
       select Finish partitioning and write changes to disk
d-i partman/confirm boolean true
```

B.4.6. Clock and time zone setup

```
# Controls whether or not the hardware clock is set to UTC.
d-i clock-setup/utc boolean true

# You may set this to any valid setting for $TZ; see the contents of
# /usr/share/zoneinfo/ for valid values.
d-i time/zone string US/Eastern
```

B.4.7. Apt setup

Setup of the `/etc/apt/sources.list` and basic configuration options is fully automated based on your installation method and answers to earlier questions. You can optionally add other (local) repositories.

```
# You can choose to install non-free and contrib software.
#d-i apt-setup/non-free boolean true
#d-i apt-setup/contrib boolean true
# Uncomment this if you don't want to use a network mirror.
#d-i apt-setup/use_mirror boolean false
# Uncomment this to avoid adding security sources, or
# add a hostname to use a different server than security.debian.org.
#d-i apt-setup/security_host string

# Additional repositories, local[0-9] available
#d-i apt-setup/local0/repository string \
#       deb http://local.server/debian stable main
#d-i apt-setup/local0/comment string local server
# Enable deb-src lines
#d-i apt-setup/local0/source boolean true
# URL to the public key of the local repository; you must provide a key or
# apt will complain about the unauthenticated repository and so the
# sources.list line will be left commented out
#d-i apt-setup/local0/key string http://local.server/key

# By default the installer requires that repositories be authenticated
# using a known gpg key. This setting can be used to disable that
# authentication. Warning: Insecure, not recommended.
#d-i debian-installer/allow_unauthenticated string true
```

B.4.8. Account setup

The password for the root account and name and password for a first regular user's account can be preseeded. For the passwords you can use either clear text values or MD5 *hashes*.

Warning

Be aware that preseeding passwords is not completely secure as everyone with access to the preconfiguration file will have the knowledge of these passwords. Using MD5 hashes is considered slightly better in terms of security but it might also give a false sense of security as access to a MD5 hash allows for brute force attacks.

```
# Skip creation of a root account (normal user account will be able to
# use sudo).
#d-i passwd/root-login boolean false
# Alternatively, to skip creation of a normal user account.
#d-i passwd/make-user boolean false

# Root password, either in clear text
#d-i passwd/root-password password r00tme
#d-i passwd/root-password-again password r00tme
# or encrypted using an MD5 hash.
#d-i passwd/root-password-crypted password [MD5 hash]

# To create a normal user account.
#d-i passwd/user-fullname string Debian User
#d-i passwd/username string debian
# Normal user's password, either in clear text
#d-i passwd/user-password password insecure
#d-i passwd/user-password-again password insecure
# or encrypted using an MD5 hash.
#d-i passwd/user-password-crypted password [MD5 hash]
```

The `passwd/root-password-crypted` and `passwd/user-password-crypted` variables can also be preseeded with "!" as their value. In that case, the corresponding account is disabled. This may be convenient for the root account, provided of course that an alternative method is setup to allow administrative activities or root login (for instance by using SSH key authentication or **sudo**).

An MD5 hash for a password can be generated using the following command.

```
$ echo "r00tme" | mkpasswd -s -H MD5
```

B.4.9. Base system installation

There is actually not very much that can be preseeded for this stage of the installation. The only questions asked concern the installation of the kernel.

```
# Select the initramfs generator used to generate the initrd for 2.6 kernels.
#d-i base-installer/kernel/linux/initramfs-generators string yaird
```

B.4.10. Boot loader installation

```
# Grub is the default boot loader (for x86). If you want lilo installed
# instead, uncomment this:
#d-i grub-installer/skip boolean true

# This is fairly safe to set, it makes grub install automatically to the MBR
# if no other operating system is detected on the machine.
d-i grub-installer/only_debian boolean true
```

```
# This one makes grub-installer install to the MBR if it also finds some other
# OS, which is less safe as it might not be able to boot that other OS.
d-i grub-installer/with_other_os boolean true

# Alternatively, if you want to install to a location other than the mbr,
# uncomment and edit these lines:
#d-i grub-installer/only_debian boolean false
#d-i grub-installer/with_other_os boolean false
#d-i grub-installer/bootdev  string (hd0,0)
# To install grub to multiple disks:
#d-i grub-installer/bootdev  string (hd0,0) (hd1,0) (hd2,0)
```

B.4.11. Package selection

You can choose to install any combination of tasks that are available. Available tasks as of this writing include:

- **standard**
- **desktop**
- **gnome-desktop**
- **kde-desktop**
- **web-server**
- **print-server**
- **dns-server**
- **file-server**
- **mail-server**
- **sql-database**
- **laptop**

You can also choose to install no tasks, and force the installation of a set of packages in some other way. We recommend always including the **standard** task.

If you want to install some individual packages in addition to packages installed by tasks, you can use the parameter `pkgsel/include`. The value of this parameter can be a list of packages separated by either commas or spaces, which allows it to be used easily on the kernel command line as well.

```
tasksel tasksel/first multiselect standard, desktop
#tasksel tasksel/first multiselect standard, web-server
#tasksel tasksel/first multiselect standard, kde-desktop

# Individual additional packages to install
#d-i pkgsel/include string openssh-server build-essential

# Some versions of the installer can report back on what software you have
# installed, and what software you use. The default is not to report back,
# but sending reports helps the project determine what software is most
# popular and include it on CDs.
#popularity-contest popularity-contest/participate boolean false
```

B.4.12. Finishing up the first stage install

```
# Avoid that last message about the install being complete.
d-i finish-install/reboot_in_progress note

# This will prevent the installer from ejecting the CD during the reboot,
# which is useful in some situations.
#d-i cdrom-detect/eject boolean false
```

B.4.13. X configuration

Preseeding Debian's X config is possible, but you probably need to know some details about the video hardware of the machine, since Debian's X configurator does not do fully automatic configuration of everything.

```
# X can detect the right driver for some cards, but if you're preseeding,
# you override whatever it chooses. Still, vesa will work most places.
#xserver-xorg xserver-xorg/config/device/driver select vesa

# A caveat with mouse autodetection is that if it fails, X will retry it
# over and over. So if it's preseeded to be done, there is a possibility of
# an infinite loop if the mouse is not autodetected.
#xserver-xorg xserver-xorg/autodetect_mouse boolean true

# Monitor autodetection is recommended.
xserver-xorg xserver-xorg/autodetect_monitor boolean true
# Uncomment if you have an LCD display.
#xserver-xorg xserver-xorg/config/monitor/lcd boolean true
# X has three configuration paths for the monitor. Here's how to preseed
# the "medium" path, which is always available. The "simple" path may not
# be available, and the "advanced" path asks too many questions.
xserver-xorg xserver-xorg/config/monitor/selection-method \
       select medium
xserver-xorg xserver-xorg/config/monitor/mode-list \
       select 1024x768 @ 60 Hz
```

B.4.14. Preseeding other packages

```
# Depending on what software you choose to install, or if things go wrong
# during the installation process, it's possible that other questions may
# be asked. You can preseed those too, of course. To get a list of every
# possible question that could be asked during an install, do an
# installation, and then run these commands:
#   debconf-get-selections --installer > file
#   debconf-get-selections >> file
```

B.5. Advanced options

B.5.1. Running custom commands during the installation

A very powerful and flexible option offered by the preconfiguration tools is the ability to run commands or scripts at certain points in the installation.

```
# d-i preseeding is inherently not secure. Nothing in the installer checks
# for attempts at buffer overflows or other exploits of the values of a
# preconfiguration file like this one. Only use preconfiguration files from
# trusted locations! To drive that home, and because it's generally useful,
# here's a way to run any shell command you'd like inside the installer,
# automatically.

# This first command is run as early as possible, just after
# preseeding is read.
#d-i preseed/early_command string anna-install some-udeb

# This command is run just before the install finishes, but when there is
# still a usable /target directory. You can chroot to /target and use it
# directly, or use the apt-install and in-target commands to easily install
# packages and run commands in the target system.
#d-i preseed/late_command string apt-install zsh; in-target chsh -s /bin/zsh
```

B.5.2. Using preseeding to change default values

It is possible to use preseeding to change the default answer for a question, but still have the question asked. To do this the *seen* flag must be reset to "false" after setting the value for a question.

```
d-i foo/bar string value
d-i foo/bar seen false
```

The same effect can be achieved for *all* questions by setting the parameter `preseed/interactive=true` at the boot prompt. This can also be useful for testing or debugging your preconfiguration file. If you are preseeding using boot parameters, you can make the installer ask the corresponding question by using the "?=" operator, i.e. ***foo/bar?=value***. This will of course only have effect for parameters that correspond to questions that are actually displayed during an installation and not for "internal" parameters.

B.5.3. Chainloading preconfiguration files

It is possible to include other preconfiguration files from a preconfiguration file. Any settings in those files will override pre-existing settings from files loaded earlier. This makes it possible to put, for example, general networking settings for your location in one file and more specific settings for certain configurations in other files.

```
# More than one file can be listed, separated by spaces; all will be
# loaded. The included files can have preseed/include directives of their
# own as well. Note that if the filenames are relative, they are taken from
# the same directory as the preconfiguration file that includes them.
```

```
#d-i preseed/include string x.cfg

# The installer can optionally verify checksums of preconfiguration files
# before using them. Currently only md5sums are supported, list the md5sums
# in the same order as the list of files to include.
#d-i preseed/include/checksum string 5da499872becccfeda2c4872f9171c3d

# More flexibly, this runs a shell command and if it outputs the names of
# preconfiguration files, includes those files.
#d-i preseed/include_command \
#      string echo if [ "‘hostname‘" = bob ]; then echo bob.cfg; fi

# Most flexibly of all, this downloads a program and runs it. The program
# can use commands such as debconf-set to manipulate the debconf database.
# More than one script can be listed, separated by spaces.
# Note that if the filenames are relative, they are taken from the same
# directory as the preconfiguration file that runs them.
#d-i preseed/run string foo.sh
```

It is also possible to chainload from the initrd or file preseeding phase, into network preseeding by setting preseed/url in the earlier files. This will cause network preseeding to be performed when the network comes up. You need to be careful when doing this, since there will be two distinct runs at preseeding, meaning for example that you get another chance to run the preseed/early command, the second one happening after the network comes up.

Appendix C. Partitioning for Debian

C.1. Deciding on Debian Partitions and Sizes

At a bare minimum, GNU/Linux needs one partition for itself. You can have a single partition containing the entire operating system, applications, and your personal files. Most people feel that a separate swap partition is also a necessity, although it's not strictly true. "Swap" is scratch space for an operating system, which allows the system to use disk storage as "virtual memory". By putting swap on a separate partition, Linux can make much more efficient use of it. It is possible to force Linux to use a regular file as swap, but it is not recommended.

Most people choose to give GNU/Linux more than the minimum number of partitions, however. There are two reasons you might want to break up the file system into a number of smaller partitions. The first is for safety. If something happens to corrupt the file system, generally only one partition is affected. Thus, you only have to replace (from the backups you've been carefully keeping) a portion of your system. At a bare minimum, you should consider creating what is commonly called a "root partition". This contains the most essential components of the system. If any other partitions get corrupted, you can still boot into GNU/Linux to fix the system. This can save you the trouble of having to reinstall the system from scratch.

The second reason is generally more important in a business setting, but it really depends on your use of the machine. For example, a mail server getting spammed with e-mail can easily fill a partition. If you made `/var/mail` a separate partition on the mail server, most of the system will remain working even if you get spammed.

The only real drawback to using more partitions is that it is often difficult to know in advance what your needs will be. If you make a partition too small then you will either have to reinstall the system or you will be constantly moving things around to make room in the undersized partition. On the other hand, if you make the partition too big, you will be wasting space that could be used elsewhere. Disk space is cheap nowadays, but why throw your money away?

C.2. The Directory Tree

Debian GNU/Linux adheres to the Filesystem Hierarchy Standard (http://www.pathname.com/fhs/) for directory and file naming. This standard allows users and software programs to predict the location of files and directories. The root level directory is represented simply by the slash `/`. At the root level, all Debian systems include these directories:

Directory	Content
`bin`	Essential command binaries
`boot`	Static files of the boot loader
`dev`	Device files
`etc`	Host-specific system configuration
`home`	User home directories
`lib`	Essential shared libraries and kernel modules
`media`	Contains mount points for replaceable media

Directory	Content
mnt	Mount point for mounting a file system temporarily
proc	Virtual directory for system information (2.4 and 2.6 kernels)
root	Home directory for the root user
sbin	Essential system binaries
sys	Virtual directory for system information (2.6 kernels)
tmp	Temporary files
usr	Secondary hierarchy
var	Variable data
srv	Data for services provided by the system
opt	Add-on application software packages

The following is a list of important considerations regarding directories and partitions. Note that disk usage varies widely given system configuration and specific usage patterns. The recommendations here are general guidelines and provide a starting point for partitioning.

- The root partition `/` must always physically contain `/etc`, `/bin`, `/sbin`, `/lib` and `/dev`, otherwise you won't be able to boot. Typically 150–250MB is needed for the root partition.
- `/usr`: contains all user programs (`/usr/bin`), libraries (`/usr/lib`), documentation (`/usr/share/doc`), etc. This is the part of the file system that generally takes up most space. You should provide at least 500MB of disk space. This amount should be increased depending on the number and type of packages you plan to install. A generous workstation or server installation should allow 4–6GB.
- `/var`: variable data like news articles, e-mails, web sites, databases, the packaging system cache, etc. will be placed under this directory. The size of this directory depends greatly on the usage of your system, but for most people will be dictated by the package management tool's overhead. If you are going to do a full installation of just about everything Debian has to offer, all in one session, setting aside 2 or 3 GB of space for `/var` should be sufficient. If you are going to install in pieces (that is to say, install services and utilities, followed by text stuff, then X, ...), you can get away with 300–500 MB. If hard drive space is at a premium and you don't plan on doing major system updates, you can get by with as little as 30 or 40 MB.
- `/tmp`: temporary data created by programs will most likely go in this directory. 40–100MB should usually be enough. Some applications — including archive manipulators, CD/DVD authoring tools, and multimedia software — may use `/tmp` to temporarily store image files. If you plan to use such applications, you should adjust the space available in `/tmp` accordingly.
- `/home`: every user will put his personal data into a subdirectory of this directory. Its size depends on how many users will be using the system and what files are to be stored in their directories. Depending on your planned usage you should reserve about 100MB for each user, but adapt this value to your needs. Reserve a lot more space if you plan to save a lot of multimedia files (pictures, MP3, movies) in your home directory.

C.3. Recommended Partitioning Scheme

For new users, personal Debian boxes, home systems, and other single-user setups, a single `/` partition (plus swap) is probably the easiest, simplest way to go. However, if your partition is larger than around 6GB, choose ext3 as your partition type. Ext2 partitions need periodic file system integrity checking, and this can cause delays during booting when the partition is large.

For multi-user systems or systems with lots of disk space, it's best to put `/usr`, `/var`, `/tmp`, and `/home` each on their own partitions separate from the `/` partition.

You might need a separate `/usr/local` partition if you plan to install many programs that are not part of the Debian distribution. If your machine will be a mail server, you might need to make `/var/mail` a separate partition. Often, putting `/tmp` on its own partition, for instance 20–50MB, is a good idea. If you are setting up a server with lots of user accounts, it's generally good to have a separate, large `/home` partition. In general, the partitioning situation varies from computer to computer depending on its uses.

For very complex systems, you should see the Multi Disk HOWTO (http://www.tldp.org/HOWTO/Multi-Disk-HOWTO.html). This contains in-depth information, mostly of interest to ISPs and people setting up servers.

With respect to the issue of swap partition size, there are many views. One rule of thumb which works well is to use as much swap as you have system memory. It also shouldn't be smaller than 16MB, in most cases. Of course, there are exceptions to these rules. If you are trying to solve 10000 simultaneous equations on a machine with 256MB of memory, you may need a gigabyte (or more) of swap.

On 32-bit architectures (i386, m68k, 32-bit SPARC, and PowerPC), the maximum size of a swap partition is 2GB. That should be enough for nearly any installation. However, if your swap requirements are this high, you should probably try to spread the swap across different disks (also called "spindles") and, if possible, different SCSI or IDE channels. The kernel will balance swap usage between multiple swap partitions, giving better performance.

As an example, an older home machine might have 32MB of RAM and a 1.7GB IDE drive on `/dev/hda`. There might be a 500MB partition for another operating system on `/dev/hda1`, a 32MB swap partition on `/dev/hda3` and about 1.2GB on `/dev/hda2` as the Linux partition.

For an idea of the space taken by tasks you might be interested in adding after your system installation is complete, check Section D.2.

C.4. Device Names in Linux

Linux disks and partition names may be different from other operating systems. You need to know the names that Linux uses when you create and mount partitions. Here's the basic naming scheme:

- The first floppy drive is named `/dev/fd0`.
- The second floppy drive is named `/dev/fd1`.
- The first SCSI disk (SCSI ID address-wise) is named `/dev/sda`.
- The second SCSI disk (address-wise) is named `/dev/sdb`, and so on.
- The first SCSI CD-ROM is named `/dev/scd0`, also known as `/dev/sr0`.
- The master disk on IDE primary controller is named `/dev/hda`.
- The slave disk on IDE primary controller is named `/dev/hdb`.

- The master and slave disks of the secondary controller can be called `/dev/hdc` and `/dev/hdd`, respectively. Newer IDE controllers can actually have two channels, effectively acting like two controllers.

The partitions on each disk are represented by appending a decimal number to the disk name: `sda1` and `sda2` represent the first and second partitions of the first SCSI disk drive in your system.

Here is a real-life example. Let's assume you have a system with 2 SCSI disks, one at SCSI address 2 and the other at SCSI address 4. The first disk (at address 2) is then named `sda`, and the second `sdb`. If the `sda` drive has 3 partitions on it, these will be named `sda1`, `sda2`, and `sda3`. The same applies to the `sdb` disk and its partitions.

Note that if you have two SCSI host bus adapters (i.e., controllers), the order of the drives can get confusing. The best solution in this case is to watch the boot messages, assuming you know the drive models and/or capacities.

C.5. Debian Partitioning Programs

Several varieties of partitioning programs have been adapted by Debian developers to work on various types of hard disks and computer architectures. Following is a list of the program(s) applicable for your architecture.

partman

Recommended partitioning tool in Debian. This Swiss army knife can also resize partitions, create filesystems and assign them to the mountpoints.

cfdisk

A simple-to-use, full-screen disk partitioner for the rest of us.

Note that **cfdisk** doesn't understand FreeBSD partitions at all, and, again, device names may differ as a result.

One of these programs will be run by default when you select Partition disks (or similar). It may be possible to use a different partitioning tool from the command line on VT2, but this is not recommended.

C.5.1. Partitioning for PA-RISC

PALO, the HPPA boot loader, requires a partition of type "F0" somewhere in the first 2GB. This is where the boot loader and an optional kernel and RAMdisk will be stored, so make it big enough for that — at least 4Mb (I like 8–16MB). An additional requirement of the firmware is that the Linux kernel must reside within the first 2GB of the disk. This is typically achieved by making the root ext2 partition fit entirely within the first 2GB of the disk. Alternatively you can create a small ext2 partition near the start of the disk and mount that on `/boot`, since that is the directory where the Linux kernel(s) will be stored. `/boot` needs to be big enough to hold whatever kernels (and backups) you might wish to load; 25–50MB is generally sufficient.

Appendix D. Random Bits

D.1. Linux Devices

In Linux various special files can be found under the directory `/dev`. These files are called device files and behave unlike ordinary files. The most common types of device files are for block devices and character devices. These files are an interface to the actual driver (part of the Linux kernel) which in turn accesses the hardware. Another, less common, type of device file is the named *pipe*. The most important device files are listed in the tables below.

`fd0`	First Floppy Drive
`fd1`	Second Floppy Drive

`hda`	IDE Hard disk / CD-ROM on the first IDE port (Master)
`hdb`	IDE Hard disk / CD-ROM on the first IDE port (Slave)
`hdc`	IDE Hard disk / CD-ROM on the second IDE port (Master)
`hdd`	IDE Hard disk / CD-ROM on the second IDE port (Slave)
`hda1`	First partition of the first IDE hard disk
`hdd15`	Fifteenth partition of the fourth IDE hard disk

`sda`	SCSI Hard disk with lowest SCSI ID (e.g. 0)
`sdb`	SCSI Hard disk with next higher SCSI ID (e.g. 1)
`sdc`	SCSI Hard disk with next higher SCSI ID (e.g. 2)
`sda1`	First partition of the first SCSI hard disk
`sdd10`	Tenth partition of the fourth SCSI hard disk

`sr0`	SCSI CD-ROM with the lowest SCSI ID
`sr1`	SCSI CD-ROM with the next higher SCSI ID

`ttyS0`	Serial port 0, COM1 under MS-DOS
`ttyS1`	Serial port 1, COM2 under MS-DOS
`psaux`	PS/2 mouse device
`gpmdata`	Pseudo device, repeater data from GPM (mouse) daemon

`cdrom`	Symbolic link to the CD-ROM drive

`mouse`	Symbolic link to the mouse device file

`null`	Anything written to this device will disappear
`zero`	One can endlessly read zeros out of this device

D.1.1. Setting Up Your Mouse

The mouse can be used in both the Linux console (with gpm) and the X window environment. Normally, this is a simple matter of installing `gpm` and the X server itself. Both should be configured to use `/dev/input/mice` as the mouse device. The correct mouse protocol is named **`exps2`** in gpm, and **`ExplorerPS/2`** in X. The respective configuration files are `/etc/gpm.conf` and `/etc/X11/xorg.conf`.

Certain kernel modules must be loaded in order for your mouse to work. In most cases the correct modules are autodetected, but not always for old-style serial and bus mice[1], which are quite rare except on very old computers. Summary of Linux kernel modules needed for different mouse types:

Module	Description
psmouse	PS/2 mice (should be autodetected)
usbhid	USB mice (should be autodetected)
sermouse	Most serial mice
logibm	Bus mouse connected to Logitech adapter card
inport	Bus mouse connected to ATI or Microsoft InPort card

To load a mouse driver module, you can use the **modconf** command (from the package with the same name) and look in the category **`kernel/drivers/input/mouse`**.

D.2. Disk Space Needed for Tasks

A standard installation for the i386 architecture, including all standard packages and using the default 2.6 kernel, takes up 353MB of disk space. A minimal base installation, without the “Standard system” task selected, will take 225MB.

> **Important:** In both cases this is the actual disk space used *after* the installation is finished and any temporary files deleted. It also does not take into account overhead used by the file system, for example for journal files. This means that significantly more disk space is needed both *during* the installation and for normal system use.

The following table lists sizes reported by aptitude for the tasks listed in tasksel. Note that some tasks have overlapping constituents, so the total installed size for two tasks together may be less than the

1. Serial mice usually have a 9-hole D-shaped connector; bus mice have an 8-pin round connector, not to be confused with the 6-pin round connector of a PS/2 mouse or the 4-pin round connector of an ADB mouse.

total obtained by adding up the numbers.

Note that you will need to add the sizes listed in the table to the size of the standard installation when determining the size of partitions. Most of the size listed as "Installed size" will end up in `/usr` and in `/lib`; the size listed as "Download size" is (temporarily) required in `/var`.

Task	Installed size (MB)	Download size (MB)	Space needed to install (MB)
Desktop environment	1360	454	1814
Laptop[a]	27	10	37
Web server	35	11	46
Print server	184	56	240
DNS server	2	1	3
File server	50	21	71
Mail server	13	5	18
SQL database	32	10	42
Notes: a. There is a large overlap of the Laptop task with the Destop environment task. If you install both, the Laptop task will only require a few MB additional disk space.			

Note: The *Desktop* task will install the GNOME desktop environment.

If you install in a language other than English, **tasksel** may automatically install a *localization task*, if one is available for your language. Space requirements differ per language; you should allow up to 350MB in total for download and installation.

D.3. Installing Debian GNU/Linux from a Unix/Linux System

This section explains how to install Debian GNU/Linux from an existing Unix or Linux system, without using the menu-driven installer as explained in the rest of the manual. This "cross-install" HOWTO has been requested by users switching to Debian GNU/Linux from Red Hat, Mandrake, and SUSE. In this section some familiarity with entering *nix commands and navigating the file system is assumed. In this section, `$` symbolizes a command to be entered in the user's current system, while # refers to a command entered in the Debian chroot.

Once you've got the new Debian system configured to your preference, you can migrate your existing user data (if any) to it, and keep on rolling. This is therefore a "zero downtime" Debian GNU/Linux install. It's also a clever way for dealing with hardware that otherwise doesn't play friendly with various boot or installation media.

D.3.1. Getting Started

With your current *nix partitioning tools, repartition the hard drive as needed, creating at least one filesystem plus swap. You need around 350MB of space available for a console only install, or about

1GB if you plan to install X (more if you intend to install desktop environments like GNOME or KDE).

Next, create file systems on the partitions. For example, to create an ext3 file system on partition `/dev/hda6` (that's our example root partition):

```
# mke2fs -j /dev/hda6
```

To create an ext2 file system instead, omit **`-j`**.

Initialize and activate swap (substitute the partition number for your intended Debian swap partition):

```
# mkswap /dev/hda5
# sync; sync; sync
# swapon /dev/hda5
```

Mount one partition as `/mnt/debinst` (the installation point, to be the root (`/`) filesystem on your new system). The mount point name is strictly arbitrary, it is referenced later below.

```
# mkdir /mnt/debinst
# mount /dev/hda6 /mnt/debinst
```

> **Note:** If you want to have parts of the filesystem (e.g. /usr) mounted on separate partitions, you will need to create and mount these directories manually before proceding with the next stage.

D.3.2. Install debootstrap

The utility used by the Debian installer, and recognized as the official way to install a Debian base system, is **debootstrap**. It uses **wget** and **ar**, but otherwise depends only on `/bin/sh` and basic Unix/Linux tools[2]. Install **wget** and **ar** if they aren't already on your current system, then download and install **debootstrap**.

Or, you can use the following procedure to install it manually. Make a work folder for extracting the .deb into:

```
# mkdir work
# cd work
```

The **debootstrap** binary is located in the Debian archive (be sure to select the proper file for your architecture). Download the **debootstrap** .deb from the pool (http://ftp.debian.org/debian/pool/main/d/debootstrap/), copy the package to the work folder, and extract the files from it. You will need to have root privileges to install the files.

```
# ar -x debootstrap_0.X.X_all.deb
# cd /
# zcat /full-path-to-work/work/data.tar.gz | tar xv
```

2. These include the GNU core utilities and commands like **sed**, **grep**, **tar** and **gzip**.

D.3.3. Run debootstrap

debootstrap can download the needed files directly from the archive when you run it. You can substitute any Debian archive mirror for **`http.us.debian.org/debian`** in the command example below, preferably a mirror close to you network-wise. Mirrors are listed at http://www.debian.org/misc/README.mirrors.

If you have a etch Debian GNU/Linux CD mounted at `/cdrom`, you could substitute a file URL instead of the http URL: **`file:/cdrom/debian/`**

Substitute one of the following for *`ARCH`* in the **debootstrap** command: **`alpha`**, **`amd64`**, **`arm`**, **`hppa`**, **`i386`**, **`ia64`**, **`m68k`**, **`mips`**, **`mipsel`**, **`powerpc`**, **`s390`**, or **`sparc`**.

```
# /usr/sbin/debootstrap --arch ARCH etch \
     /mnt/debinst http://ftp.us.debian.org/debian
```

D.3.4. Configure The Base System

Now you've got a real Debian system, though rather lean, on disk. **chroot** into it:

```
# LANG=C chroot /mnt/debinst /bin/bash
```

After chrooting you may need to set the terminal definition to be compatible with the Debian base system, for example:

```
# export TERM=xterm-color
```

D.3.4.1. Mount Partitions

You need to create `/etc/fstab`.

```
# editor /etc/fstab
```

Here is a sample you can modify to suit:

```
# /etc/fstab: static file system information.
#
# file system    mount point   type    options                 dump pass
/dev/XXX         /             ext3    defaults                0    1
/dev/XXX         /boot         ext3    ro,nosuid,nodev         0    2

/dev/XXX         none          swap    sw                      0    0
proc             /proc         proc    defaults                0    0

/dev/fd0         /media/floppy auto    noauto,rw,sync,user,exec 0    0
```

```
/dev/cdrom       /media/cdrom  iso9660 noauto,ro,user,exec      0    0

/dev/XXX         /tmp          ext3    rw,nosuid,nodev          0    2
/dev/XXX         /var          ext3    rw,nosuid,nodev          0    2
/dev/XXX         /usr          ext3    rw,nodev                 0    2
/dev/XXX         /home         ext3    rw,nosuid,nodev          0    2
```

Use **mount -a** to mount all the file systems you have specified in your /etc/fstab, or, to mount file systems individually, use:

```
# mount /path   # e.g.: mount /usr
```

Current Debian systems have mountpoints for removable media under /media, but keep compatibility symlinks in /. Create these as as needed, for example:

```
# cd /media
# mkdir cdrom0
# ln -s cdrom0 cdrom
# cd /
# ln -s media/cdrom
```

You can mount the proc file system multiple times and to arbitrary locations, though /proc is customary. If you didn't use **mount -a**, be sure to mount proc before continuing:

```
# mount -t proc proc /proc
```

The command **ls /proc** should now show a non-empty directory. Should this fail, you may be able to mount proc from outside the chroot:

```
# mount -t proc proc /mnt/debinst/proc
```

D.3.4.2. Setting Timezone

An option in the file /etc/default/rcS determines whether the system will interpret the hardware clock as being set to UTC or local time. The following command allow you to set that and choose your timezone.

```
# editor /etc/default/rcS
# tzconfig
```

D.3.4.3. Configure Networking

To configure networking, edit /etc/network/interfaces, /etc/resolv.conf, /etc/hostname and /etc/hosts.

```
# editor /etc/network/interfaces
```

Here are some simple examples from `/usr/share/doc/ifupdown/examples`:

```
#######################################################################
# /etc/network/interfaces -- configuration file for ifup(8), ifdown(8)
# See the interfaces(5) manpage for information on what options are
# available.
#######################################################################

# We always want the loopback interface.
#
auto lo
iface lo inet loopback

# To use dhcp:
#
# auto eth0
# iface eth0 inet dhcp

# An example static IP setup: (broadcast and gateway are optional)
#
# auto eth0
# iface eth0 inet static
#     address 192.168.0.42
#     network 192.168.0.0
#     netmask 255.255.255.0
#     broadcast 192.168.0.255
#     gateway 192.168.0.1
```

Enter your nameserver(s) and search directives in `/etc/resolv.conf`:

```
# editor /etc/resolv.conf
```

A simple example `/etc/resolv.conf`:

```
search hqdom.local
nameserver 10.1.1.36
nameserver 192.168.9.100
```

Enter your system's host name (2 to 63 characters):

```
# echo DebianHostName > /etc/hostname
```

And a basic `/etc/hosts` with IPv6 support:

```
127.0.0.1 localhost DebianHostName

# The following lines are desirable for IPv6 capable hosts
::1     ip6-localhost ip6-loopback
fe00::0 ip6-localnet
ff00::0 ip6-mcastprefix
ff02::1 ip6-allnodes
ff02::2 ip6-allrouters
ff02::3 ip6-allhosts
```

If you have multiple network cards, you should arrange the names of driver modules in the `/etc/modules` file into the desired order. Then during boot, each card will be associated with the interface name (eth0, eth1, etc.) that you expect.

D.3.4.4. Configure Apt

Debootstrap will have created a very basic `/etc/apt/sources.list` that will allow installing additional packages. However, you may want to add some additional sources, for example for source packages and security updates:

```
deb-src http://ftp.us.debian.org/debian etch main

deb http://security.debian.org/ etch/updates main
deb-src http://security.debian.org/ etch/updates main
```

Make sure to run **`aptitude update`** after you have made changes to the sources list.

D.3.4.5. Configure Locales and Keyboard

To configure your locale settings to use a language other than English, install the `locales` support package and configure it. Currently the use of UTF-8 locales is recommended.

```
# aptitude install locales
# dpkg-reconfigure locales
```

To configure your keyboard (if needed):

```
# aptitude install console-data
# dpkg-reconfigure console-data
```

Note that the keyboard cannot be set while in the chroot, but will be configured for the next reboot.

D.3.5. Install a Kernel

If you intend to boot this system, you probably want a Linux kernel and a boot loader. Identify available pre-packaged kernels with:

```
# apt-cache search linux-image
```

If you intend to use a pre-packaged kernel, you may want to create the configuration file `/etc/kernel-img.conf` before you do so. Here's an example file:

```
# Kernel image management overrides
# See kernel-img.conf(5) for details
do_symlinks = yes
relative_links = yes
do_bootloader = yes
```

```
do_bootfloppy = no
do_initrd = yes
link_in_boot = no
```

For detailed information about this file and the various options, consult its man page which will be available after installing the `kernel-package` package. We recommend that you check that the values are appropriate for your system.

Then install the kernel package of your choice using its package name.

```
# aptitude install linux-image-2.6.18-arch-etc
```

If you did not create a `/etc/kernel-img.conf` before installing a pre-packaged kernel, you may be asked some questions during its installation that refer to it.

D.3.6. Set up the Boot Loader

To make your Debian GNU/Linux system bootable, set up your boot loader to load the installed kernel with your new root partition. Note that **debootstrap** does not install a boot loader, though you can use **aptitude** inside your Debian chroot to do so.

D.3.7. Finishing touches

As mentioned earlier, the installed system will be very basic. If you would like to make the system a bit more mature, there is an easy method to install all packages with “standard” priority:

```
# tasksel install standard
```

Of course, you can also just use **aptitude** to install packages individually.

After the installation there will be a lot of downloaded packages in `/var/cache/apt/archives/`. You can free up some diskspace by running:

```
# aptitude clean
```

D.4. Installing Debian GNU/Linux using PPP over Ethernet (PPPoE)

In some countries PPP over Ethernet (PPPoE) is a common protocol for broadband (ADSL or cable) connections to an Internet Service Provider. Setting up a network connection using PPPoE is not supported by default in the installer, but can be made to work very simply. This section explains how.

The PPPoE connection set up during the installation will also be available after the reboot into the installed system (see Chapter 7).

To have the option of setting up and using PPPoE during the installation, you will need to install using one of the CD-ROM/DVD images that are available. It is not supported for other installation methods (e.g. netboot).

Installing over PPPoE is mostly the same as any other installation. The following steps explain the differences.

- Boot the installer with **`modules=ppp-udeb`** as boot parameter. This means that at the boot prompt you should enter:

  ```
  install modules=ppp-udeb
  ```

  or, if you prefer using the graphical installer:

  ```
  installgui modules=ppp-udeb
  ```

  This will ensure the component responsible for the setup of PPPoE (`ppp-udeb`) will be loaded and run automatically.
- Follow the regular initial steps of the installation (language, country and keyboard selection; the loading of additional installer components[3]).
- The next step is the detection of network hardware, in order to identify any Ethernet cards present in the system.
- After this the actual setup of PPPoE is started. The installer will probe all the detected Ethernet interfaces in an attempt to find a PPPoE concentrator (a type of server which handles PPPoE connections).

  It is possible that the concentrator will not to be found at the first attempt. This can happen occasionally on slow or loaded networks or with faulty servers. In most cases a second attempt to detect the concentrator will be successful; to retry, select Configure and start a PPPoE connection from the main menu of the installer.
- After a concentrator is found, the user will be prompted to type the login information (the PPPoE username and password).
- At this point the installer will use the provided information to establish the PPPoE connection. If the correct information was provided, the PPPoE connection should be configured and the installer should be able to use it to connect to the Internet and retrieve packages over it (if needed). If the login information is not correct or some error appears, the installer will stop, but the configuration can be attempted again by selecting the menu entry Configure and start a PPPoE connection.

3. The `ppp-udeb` component is loaded as one of the additional components in this step. If you want to install at medium or low priority (expert mode), you can also manually select the `ppp-udeb` instead of entering the “modules” parameter at the boot prompt.

Appendix E. Administrivia

E.1. About This Document

This manual was created for Sarge's debian-installer, based on the Woody installation manual for boot-floppies, which was based on earlier Debian installation manuals, and on the Progeny distribution manual which was released under GPL in 2003.

This document is written in DocBook XML. Output formats are generated by various programs using information from the `docbook-xml` and `docbook-xsl` packages.

In order to increase the maintainability of this document, we use a number of XML features, such as entities and profiling attributes. These play a role akin to variables and conditionals in programming languages. The XML source to this document contains information for each different architecture — profiling attributes are used to isolate certain bits of text as architecture-specific.

E.2. Contributing to This Document

If you have problems or suggestions regarding this document, you should probably submit them as a bug report against the package `installation-guide`. See the `reportbug` package or read the online documentation of the Debian Bug Tracking System (http://bugs.debian.org/). It would be nice if you could check the open bugs against installation-guide (http://bugs.debian.org/installation-guide) to see whether your problem has already been reported. If so, you can supply additional corroboration or helpful information to <*xxxx*`@bugs.debian.org`>, where *xxxx* is the number for the already-reported bug.

Better yet, get a copy of the DocBook source for this document, and produce patches against it. The DocBook source can be found at the debian-installer WebSVN (http://svn.debian.org/wsvn/d-i/). If you're not familiar with DocBook, don't worry: there is a simple cheatsheet in the manuals directory that will get you started. It's like html, but oriented towards the meaning of the text rather than the presentation. Patches submitted to the debian-boot mailing list (see below) are welcomed. For instructions on how to check out the sources via SVN, see README (http://svn.debian.org/wsvn/d-i/README?op=file) from the source root directory.

Please do *not* contact the authors of this document directly. There is also a discussion list for `debian-installer`, which includes discussions of this manual. The mailing list is <`debian-boot@lists.debian.org`>. Instructions for subscribing to this list can be found at the Debian Mailing List Subscription (http://www.debian.org/MailingLists/subscribe) page; or you can browse the Debian Mailing List Archives (http://lists.debian.org/) online.

E.3. Major Contributions

This document was originally written by Bruce Perens, Sven Rudolph, Igor Grobman, James Treacy, and Adam Di Carlo. Sebastian Ley wrote the Installation Howto. Many, many Debian users and developers contributed to this document. Particular note must be made of Michael Schmitz (m68k support), Frank Neumann (original author of the Amiga install manual (http://www.informatik.uni-oldenburg.de/~amigo/debian_inst.html)), Arto Astala, Eric Delaunay/Ben Collins (SPARC information), Tapio Lehtonen, and Stéphane Bortzmeyer for numerous edits and text. We have to thank Pascal

Le Bail for useful information about booting from USB memory sticks. Miroslav Kuře has documented a lot of the new functionality in Sarge's debian-installer.

Extremely helpful text and information was found in Jim Mintha's HOWTO for network booting (no URL available), the Debian FAQ (http://www.debian.org/doc/FAQ/), the Linux/m68k FAQ (http://www.linux-m68k.org/faq/faq.html), the Linux for SPARC Processors FAQ (http://www.ultralinux.org/faq.html), the Linux/Alpha FAQ (http://linux.iol.unh.edu/linux/alpha/faq/), amongst others. The maintainers of these freely available and rich sources of information must be recognized.

The section on chrooted installations in this manual (Section D.3) was derived in part from documents copyright Karsten M. Self.

E.4. Trademark Acknowledgement

All trademarks are property of their respective trademark owners.

Appendix F. GNU General Public License

Version 2, June 1991


```
Copyright (C) 1989, 1991 Free Software Foundation, Inc.
51 Franklin St, Fifth Floor, Boston, MA 02110-1301, USA.

Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.
```


F.1. Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the gnu General Public License is intended to guarantee your freedom to share and change free software — to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the gnu Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

F.2. 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 AND 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

F.3. 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) year  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., 51 Franklin Street, Fifth Floor, Boston, MA  02110-1301, 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) year 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.
```

```
signature of 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 Lesser General Public License instead of this License.