

**Answers to Frequently asked Questions
about Linux
– FAQ –
Version 1.18**

The Linux Community

June 7, 1993

Abstract

Hi Linuxers!

The original FAQ 1st version was posted on Dec. 19, 1991 by Robert Blum.

Most credits of this work to Linus, Robert and Ted. The first X11 section was written by Peter Hawkins, the rest was either on the list posted by many (real) activists, not me ;-), either in some other news groups, or else by direct posting to me (thanks Humberto, Dan, Michael, Drew, Audoin). I haven't systematically copyrighted them, so thanks to every one who participated even indirectly to this FAQ.

Since September 1992, the FAQ is co-written by:

Who	(WHAT)	E-Mail
Michael K. Johnson	(META-FAQ)	johnson@stolaf.edu
Matt Welsh	(GENERAL INFO)	mdw@tc.cornell.edu
Zane Healy	(BBS INFO)	healyzh@holonet.net
Matt Welsh	(INSTALLATION)	mdw@tc.cornell.edu
Peter MacDonald	(SLS INFO)	pmacdona@sanjuan.uvic.ca
Mark Komarinski	(DOS)	komarimf@craft.camp.clarkson.edu
Drew Eckhardt	(SCSI)	drew@cs.colorado.edu
Rick Miller	(DEVICE INFO)	rick@ee.uwm.edu
Hongjiu Lu	(GCC)	hlu@eecs.wsu.edu
Krishna Balasubramanian	(X11)	balasub@cis.ohio-state.edu
Rick Sladkey	(EMACS)	jrs@world.std.com
Philip Copeland	(NET INFO)	p_copela@csd.bristol-poly.ac.uk
Brian McCauley	(LPD)	B.A.McCauley@bham.ac.uk
Dirk Hohndel	(PROOF READER)	hohndel@informatik.uni-wuerzburg.dbp.de
Marc-Michel Corsini	(FAQ Collector)	corsini@{labri,firmin}.greco-prog.fr

If anyone is interested in participating with this FAQ, just send me a note with: your name/e-mail and the section you want to maintain.

Many of the questions could be avoided, if people had read the FAQ of the following newsgroups: `news.announce.newusers`, `comp.lang.c.gnu.emacs.help`, `comp.unix.questions`.

[The last-change-date of this posting is always "two minutes ago". :-)]

Early \LaTeX processing by Pepe Flores P. `jflores@dcc.uchile.cl`. Thanks to him. This text is processed by Martin Schulze (`Martin.Schulze@Informatik.Uni-Oldenburg.DE`) with `emTeX`. A new copy should at least be found at `ftp{.Informatik..}Uni-Oldenburg.DE` in `\pub\linux\...` as `faq.tex.tar.{z,z}`.

Martin

Contents

1	Introduction	1
1.1	Future Plan	1
1.2	Where to get the Linux FAQ?	2
2	Warnings	2
3	Linux General Information	4
3.1	Information	4
3.1.1	What is Linux?	4
3.1.2	Does Linux support	4
3.1.3	What is the current state of Linux?	5
3.1.4	Where do I get it and how do I install it?	5
3.1.5	Is Linux really difficult to install and use?	5
3.1.6	What's the best advice you can give to a Linux newbie?	5
3.1.7	How long has Linux been publicly available?	6
3.1.8	How reliable is Linux, anyway?	6
3.1.9	What is the proper pronunciation for "Linux"?	6
3.1.10	What's about the copyright of Linux?	6
3.1.11	Should I be a UNIX and/or a DOS wizard	6
3.1.12	Why can't we split <code>comp.os.linux</code> ?	6
3.2	Requirements	7
3.2.1	Does it run on my computer?	7
3.2.2	How much space will Linux take up on my hard drive?	8
3.2.3	Why the suggested 4 MB, for Linux?	8
3.2.4	How would this operate in an OS/2 environment?	8
3.2.5	Will Linux run on a PC or 286-AT? If not, why?	8
3.2.6	Will Linux run on a 386 Laptop?	8
3.3	What Linux does	8
3.3.1	Does Linux use TSS segments to provide multitasking?	8
3.3.2	If my PC runs under Linux	8
3.3.3	Does Linux do paging?	8
3.3.4	Can I have tasks spanning the full 4GB	9
3.3.5	Does the bigger program sizes mean I can run X?	9
3.4	Comparing Linux with other Unices	9
3.4.1	What are the differences to Minix?	9
3.4.2	What are the differences to 386BSD?	9
4	Linux useful Addresses	10
4.1	Linux on the Net	10
4.1.1	Where can I get Linux?	10
4.1.2	I do not have FTP access	11
4.1.3	Who is offering mail-order Linux kits?	12
4.1.4	Is there a newsgroup or mailing-list about Linux?	12
4.1.5	Where can I get my questions answered? How about bug-reports?	12
4.1.6	Could you be more explicit about the multi-channel list?	14
4.1.7	How can I join the channel XXX?	14
4.1.8	How can I leave the channel XXX?	15
4.1.9	'm not a hacker, what are the channels?	15
4.1.10	Does there exist a place where the traffic is kept?	15
4.2	Obtaining Linux from BBS's	16
4.2.1	I don't have access to FTP, how can I obtain Linux?	16

4.2.2	I got this FAQ from a local BBS	16
4.2.3	Do BBS's offer anything that the Internet does not?	16
4.2.4	What is a BBS?	16
4.2.5	How can I get a list of BBS's that carry Linux?	16
4.2.6	I have a phone number of a BBS	16
4.2.7	There is a local BBS that carries Linux.	17
4.2.8	What can I do to help Linux?	18
4.2.9	How do I read the data in the Linux BBS List?	18
4.2.10	What are some of the best BBS's to check out?	19
4.2.11	What are File Requests?	26
5	Installation and Common Problems	26
5.1	Where to start	26
5.1.1	I want to install Linux on my machine.	26
5.1.2	Is there some kind of limit?	27
5.1.3	How can I be sure I won't be writing ...?	27
5.2	Linux Packages	27
5.2.1	Does there exist a way to get Linux or more?	27
5.2.2	Where can I get these releases of Linux?	28
5.2.3	What files do I need to get for the SLS release?	28
5.2.4	Any other information about SLS?	29
5.2.5	What should I do to install the SLS release?	31
5.2.6	When installing SLS	31
5.2.7	Where do I find the current releases of programs?	32
5.2.8	How do I install H.J. Lu's "bootable rootdisk" release?	32
5.3	Some common Problems	32
5.3.1	What filetype is the extension?	32
5.3.2	How do I make partitions for Linux?	34
5.3.3	Why does <code>fdisk</code> cannot use XXXX sectors?	36
5.3.4	What does the message from MINIX-fs mean?	36
5.3.5	Linux <code>mkfs</code> doesn't accept the size.	36
5.3.6	How can I get <code>mkfs/mkefs</code> to check for bad blocks?	36
5.3.7	Generating a bad block list for <code>mk(e)fs</code>	37
5.3.8	How can I boot Linux off of my hard drive?	40
5.3.9	I tried to install LILO, but screwed up somehow	40
5.3.10	When installing SLS, I get an error.	40
5.3.11	When installing SLS — creating a boot floppy	41
5.3.12	How can I set the default video mode?	41
5.3.13	How else can I use <code>rdev</code> ?	42
5.3.14	When I login as non-root	42
5.3.15	I have the previous Linux version, how can I upgrade it?	43
5.3.16	Why won't the kernel compile correctly?	43
5.3.17	How can I upgrade to the newest version of GCC?	44
5.3.18	What's the deal with these <i>jump tables</i> ?	44
5.3.19	How to upgrade jump tables?	45
5.3.20	How can I be sure I won't be writing over anything important?	45
5.3.21	I just rebooted my machine, and now Linux dies.	45
5.3.22	Can I use both OS/2 and Linux on my machine?	45
5.3.23	I use OS/2's Boot Manager on my hard drive	46
5.3.24	Linux's <code>fdisk</code> says OPUS for OS/2's partition.	46
5.3.25	Why doesn't the termcap entry match?	46
5.3.26	There is no <code>include/types.h</code>	46
5.3.27	What are the differences between the different system inits?	46
5.3.28	I can't get the networking to work.	46

6	Some Classical Problems	47
6.1	Kernel	47
6.1.1	While running <code>du</code> I get Kernel panic :	47
6.2	Drives	47
6.2.1	How can I partition my hard-drive to use Linux?	47
6.2.2	I heard something about repartition a hard disk	47
6.2.3	What must I do to mkfs a floppy?	47
6.2.4	I have some trouble with tar/untar .	47
6.2.5	I can do this as root but not as non-root, is it a bug?	48
6.2.6	<code>du</code> reports twice the size	48
6.2.7	Sometimes, I get mount can't open lock file	48
6.2.8	When I try <code>mount /dev/hd?? /user</code> , I get error 2.	48
6.2.9	Since I have upgrade my Linux kernel	48
6.2.10	Since Linux 0.96b I have a lot of core files.	49
6.2.11	I can only log-in as root , is it normal?	49
6.2.12	Sometimes my Linux virtual consoles get messed up	49
7	Linux and DOS	50
7.1	General	50
7.1.1	Is is possible to access DOS from Linux?	50
7.2	M-Tools	50
7.2.1	Why use mtools if you can just mount a drive?	50
7.2.2	How do I get the mtools package set up correctly?	50
7.2.3	What is the format of the file <code>/verb/etc/mtools</code> ?	50
7.2.4	Where can I find out more about mtools?	51
7.2.5	How can I install the mtools package	51
7.3	DOS filesystem	51
7.3.1	How do I use the DOS file system?	51
7.3.2	When I mount an MS-DOS disk	51
7.3.3	I want to use the DOS file system	51
7.3.4	Where can I find out more about the DOS file system	52
7.4	DOS-Emulator	52
7.4.1	What about the dos emulator (dosumu)?	52
7.4.2	I have a Stacker volume on my hard drive. Can I access that?	52
7.5	Other file systems	52
7.5.1	How about OS/2 HPFS partitions?	52
7.6	Additions	52
7.6.1	Could someone explain how to use rawrite?	52
7.6.2	What is as86.tar.Z ?	52
7.6.3	Turbo Assembler won't compile the Linux boot code.	53
8	Miscellaneous Hints	53
8.1	Misc Information	53
8.1.1	It seems that <code>#!</code> ported on Linux don't run correctly.	53
8.1.2	Has <code>#!</code> been ported to Linux?	53
8.1.3	I've ported <code>#!</code> to Linux.	54
8.1.4	I want to port <code>#!</code> to Linux, what are the flags?	54
8.2	Minor/Major Information	54
8.2.1	What are the Major/Minor device numbers	54
8.2.2	Could someone give the meaning of the IRQ's?	60
8.2.3	Could some one clear up the devices meaning?	60
8.3	Special Serial	60
8.3.1	Where can I get help with serial communications under Linux?	61
8.3.2	When I run kermit under Linux,	61

8.3.3	What are the major, minor numbers for the serial ports? . . .	61
8.3.4	A sample <code>/etc/inittab</code>	61
8.3.5	How do I set parameters like parity for serial login?	62
8.3.6	Trying to allow login on my terminal	62
8.3.7	Where are my Call-out-mode-serial ports (<code>cuax</code>)?	62
8.3.8	What other possibly free IRQ's there are than IRQ5?	62
8.3.9	Can I use IRQ 7 when not printing?	63
8.3.10	Is pin 1 nearest to the plate or is it 31?	63
9	More Hints	63
9.1	Booting	63
9.1.1	How to start Linux from drive B?	63
9.1.2	The program <code>boot_b</code> works fine.	63
9.1.3	How can I get Linux to boot directly from the harddisk? . . .	64
9.1.4	I use <code>shoelace</code>	64
9.1.5	Is it possible to disable the 3-fingers salute	64
9.1.6	Could some one explain the information provided at boot-time? 64	
9.1.7	What is the <code>rdev</code> program provided in the images?	64
9.2	Miscellaneous	65
9.2.1	How can I backup my Hd under Linux?	65
9.2.2	Where is <code>which</code> ?	65
9.2.3	How to use <code>setterm</code> : for the novice?	65
9.2.4	I've tried <code>clear/ reset</code> which exist on most of UNIX	66
9.2.5	I know there are VC, but where is the <code>setterm</code> stuff?	66
9.2.6	I know there are shared libraries	66
9.2.7	When I want to remove a directory, I get an error.	66
9.2.8	I'm looking for some stuff, where can I find them?	67
9.2.9	How can I setup a user account other than <code>root</code> ?	67
9.2.10	I've been trying to get Linux to run on my box.	67
9.2.11	Does there exist games, languages?	67
9.2.12	Whenever I use <code>uemacs 3.1X</code> on a symlink	68
9.2.13	<code>Uemacs</code> doesn't work anymore with 0.95a.	68
9.2.14	I have an SVGA, but Linux detect an EGAc/EGAm	68
9.2.15	How can I change the keyboard repeat rate?	68
9.2.16	I compiled <code>fdformat.c</code>	69
9.2.17	What is the meaning of files ended by <code>.T.Z</code> (or <code>.taz</code>)?	69
9.2.18	hat is the meaning of files ended by <code>.T.z</code> (or <code>.tpz</code>)?	69
9.2.19	I have upgraded the kernel from XX to YY.	69
9.2.20	What is <code>doshell</code> good for?	69
9.2.21	How can I change the keyboard language?	69
9.2.22	Does there exist a possibility to track down the source	70
9.2.23	How can I decode the numbers in the message?	71
10	Emacs for Linux	71
10.1	General Information	71
10.2	GNU Emacs for Linux	71
10.2.1	Does Linux Emacs use the shared libraries?	71
10.2.2	What version of the compiler was used?	71
10.2.3	Which files do I need?	71
10.2.4	How do I install them?	71
10.2.5	What if I want to compile Emacs myself?	72
10.2.6	What about Epoch or Lucid Emacs?	72
10.2.7	Does Linux Emacs support eight-bit input/output?	72
10.2.8	How much disk space is required?	72

10.2.9	Why can't Emacs find its support files anymore?	72
10.2.10	How do I get Emacs to recognize my cursor keys?	72
10.2.11	What packages are particularly useful under Linux?	72
10.2.12	Does Linux Emacs support the X Window System?	73
10.2.13	Do I need both Emacs if I don't always use X?	73
10.2.14	Why doesn't Emacs use the settings?	73
10.2.15	I read about some menu that is supposed.	73
10.2.16	Sometimes Emacs crashes	73
11	Features	73
11.1	Virtual Consoles	73
11.1.1	I've read that Linux has virtual consoles.	73
11.1.2	How can I access to the virtual consoles 5-8?	73
11.2	Miscellaneous	74
11.2.1	What kind of shell is <code>/bin/sh</code> ?	74
11.2.2	Does there exist a man page for <code>****</code> ?	74
11.2.3	Is there a simple man package (<code>groff</code> is too big)	74
11.2.4	I have a 2MB box, but <code>free</code> reports only 1MB why?	74
11.2.5	Does there exist shared libs?	74
11.2.6	Why do I need dynamic linking?	75
11.2.7	How does dynamic linking work under Linux?	76
11.3	Swapping	77
11.3.1	How do I make swapping work?	77
11.3.2	When I boot I get one of the following messages	77
11.3.3	How do I know if it is swapping?	78
11.3.4	How is it possible to remove a swap file?	78
11.3.5	How is it possible to remove a swap device?	78
11.3.6	How much swap space do I need?	78
11.3.7	Could someone explain the swap process on Linux?	78
11.3.8	Is demand paging different from paging and how?	78
11.3.9	Is there any way to tell about the used swap space ?	79
11.4	Patching	79
11.4.1	How does one go about applying a patch to Linux?	79
11.4.2	There are a lot of patches available, will they work?	79
11.4.3	I got the patches on some ftp sites, are the patches buggy?	79
11.5	Software	80
11.5.1	What are the editors available in Linux?	80
11.5.2	Does there exist a printer package for Linux?	80
11.5.3	Are USENET news available for Linux?	80
11.5.4	What is VFS?	80
11.5.5	What's about <code>TEX</code> ?	81
11.5.6	What's about <code>LILLO</code> ?	81
11.5.7	What's about <code>MGR</code> ?	81
11.5.8	I have successfully compiled <code>MGR</code> .	82
11.5.9	Any tips for <code>MGR</code> ?	82
11.5.10	What's about <code>X11</code> ?	82
11.6	Hardware	83
11.6.1	What tape drives work with Linux?	83
11.6.2	Is there only the <code>***</code> keyboard?	83
11.6.3	Some keyboard keys produce wrong characters.	83
11.6.4	Does Linux work for SCSI drives?	83
11.6.5	Linux is supposed to work with ESDI drive	83
11.6.6	What's about Bus Mice?	83

12 GCC Information	84
12.1 HLU Information	84
12.1.1 I don't know how to install gcc stuff.	84
12.1.2 What compiler should I use for Linux?	84
12.1.3 Where is the latest official gcc 2.xx for Linux?	85
12.1.4 Where is the latest official Linux C library?	85
12.1.5 What are the contents of them?	85
12.1.6 How do I install it?	85
12.1.7 What are the main differences with the old release?	85
12.1.8 Can I use the old version of gcc?	85
12.1.9 Can I delete the old shared image in /lib?	85
12.1.10 Is stdio ANSI compatible?	85
12.1.11 Is g++ in 2.xx?	85
12.1.12 Where can I get the gcc manual?	86
12.1.13 What options can I use for gcc?	86
12.1.14 How can I debug the C code?	86
12.1.15 Where is the source code of the new <code>libc.a</code> ?	86
12.1.16 Why does g++ complain, even die?	86
12.1.17 How do I generate code for i486?	86
12.1.18 I heard <code>malloc(0)</code> wouldn't work with Linux?	86
12.1.19 Why does gcc say <code>xxxxx.h not found</code> ?	86
12.1.20 I really followed every step in the documentation.	86
12.1.21 How do I compile programs under Linux?	86
12.1.22 How can I get <code>bsd</code> style signal?	87
12.1.23 Why does a program become a CPU hog?	87
12.1.24 When a program is stopped and restarted	88
12.2 Others	89
12.2.1 I seem to be unable to compile anything with gcc. Why?	89
12.2.2 gcc complains about not finding <code>crt0.o</code>	89
12.2.3 I tried to port a <i>new</i> version of gnu stuff.	90
12.2.4 How can I get <code>bsd</code> style signal?	90
12.2.5 How to compile programs which may be debugged with <code>gdb</code> ?	90
12.2.6 When compiling some code, <code>cc1</code> complains.	90
12.2.7 When compiling <code>\$#!</code> .	90
12.2.8 How can I write codes suitable for building shared library?	91
13 SCSI Special	91
13.1 General Information	91
13.1.1 What SCSI hosts are supported?	91
13.1.2 What SCSI hosts are not supported?	91
13.1.3 Which disks?	91
13.1.4 What about SCSI tapes?	92
13.1.5 What about CD ROMS?	92
13.1.6 What about flopticals, Bernoulis, Sysquests?	92
13.1.7 How do I get SCSI information?	92
13.2 Others	92
13.2.1 I can't make a filesystem on <code>/dev/hd*</code> !	92
13.2.2 How do I partition the disk?	92
13.2.3 My partitioning program can't figure out the disk geometry.	92
13.2.4 The linux partitioning programs don't work.	92
13.2.5 The stock <code>{f,pf}disk</code> programs don't work.	93
13.2.6 Where is the latest version maintained?	93
13.2.7 What are the major / minor numbers for SCSI drives?	93
13.2.8 How do I reduce kernel bloat?	93

13.2.9	I get SCSI timeouts.	93
13.2.10	The * driver doesn't work	94
13.2.11	My Seagate / Future Domain TMC-88x board is not detected.	94
13.2.12	The Seagate / Future Domain TMC-88x driver doesn't work.	94
13.2.13	When using a Seagate / Future Domain TMC-88x	95
13.2.14	The WD-7000FASST driver doesn't work	95
13.2.15	My Ultrastor 14F or 34F isn't detected	95
13.2.16	The Adaptec driver doesn't work.	95
13.2.17	My tape drive isn't recognized at boot time	95
13.2.18	My system is dog slow (ie, 60k/sec)	95
13.3	Bugs	96
13.3.1	What about bugs?	96
13.3.2	What are the known bugs?	96
13.3.3	I get a message saying READ CAPACITY FAILED	96
13.3.4	What do I do if I find a bug?	96
13.3.5	I've found one of the following bugs	97
14	X11 — the Maximum and More	97
14.1	X386 General Information	97
14.1.1	What is the X11 release supported by Linux?	97
14.1.2	What is X386/xfree86 ?	97
14.1.3	Where can I get X386 1.2 (X11R5)?	97
14.1.4	Any tips on compiling X11R5?	98
14.2	Hardware Requirements	98
14.2.1	What VGA boards are supported?	98
14.2.2	What Mouses are supported?	99
14.2.3	The <i>Mini-DIN</i> mouse on an HP Vectra 486/33T	99
14.2.4	I have trouble with my logitech Pilot mouse and X	99
14.3	Linux Distribution	100
14.3.1	What is the current version?	100
14.3.2	What Files do I need to download?	100
14.3.3	Where do the X11 files go? What are they?	100
14.4	Learning/Using X	100
14.4.1	Where can I find the basic help for learning/using X?	100
14.4.2	What docs are available besides man pages?	101
14.4.3	How do I start up X?	101
14.4.4	How do I support national keyboards in X11?	101
14.4.5	How do I configure X .. colors, menu, keyboard?	102
14.5	Debugging Startup Problems	102
14.5.1	How do I start up X?	102
14.5.2	I cannot type in my xterm	102
14.5.3	What are some of the common omissions and errors?	103
14.5.4	Why is the server unable to find some of the fonts?	103
14.5.5	My server <i>hangs</i> Why?	103
14.5.6	When I try to use X11 I got errors.	103
14.5.7	Why cant I run more than 4 xterms?	104
14.5.8	How does X11 start up?	104
14.5.9	How can I see what startx passes to?	104
14.5.10	What devices does X depend on?	104
14.6	Xconfig and Video Mode Settings	105
14.6.1	What do the mode names mean?	105
14.6.2	How do I compute the numbers for the video modes?	105
14.6.3	How can I determine the clocks on the video chip?	105
14.6.4	What are the settings used with a trident-8900C?	106

14.6.5	What can I do with a trident 8900B?	106
14.6.6	Where can I get Xconfig-files for various hardware?	106
14.6.7	What are common problems with Xconfig?	106
14.6.8	What do I use in my Xconfig file to use the bus mouse?	107
14.7	X-Applications: Compiling X Programs	107
14.7.1	Does anyone know where I can find the libobz.a library?	107
14.7.2	I get <code>_setjmp undefined</code> when linking xv from .a files.	108
14.7.3	How do I compile an X application that has an Imakefile?	108
14.7.4	How do I compile a program that has no Imakefile?	108
14.7.5	gcc complains the X libraries are not found or links static.	108
14.7.6	When using Makefiles generated by xmkmf	108
14.7.7	I recently compiled XFree86 1.2A for Linux	108
14.7.8	What are the development tools available for X under Linux?	108
14.8	SVGA server for ATI boards	109
14.8.1	What's new?	109
14.8.2	What is the ATI SVGA driver?	109
14.8.3	What is the ATI SVGA driver NOT ?	110
14.8.4	What cards will the driver work with?	110
14.8.5	What should I put in my Xconfig file?	110
14.8.6	What is the history of the driver?	112
14.8.7	What is the future of the driver?	112
14.9	Bugs	113
15	Networking and Linux	113
15.1	General	113
15.1.1	Where can I find useful information about networking for Linux? 113	
16	Everything for printing	113
16.1	General Information	114
16.1.1	What <i>is</i> the lpr/lpd package?	114
16.1.2	What is <i>the</i> Linux lpr/lpd?	114
16.1.3	What are all these directories for?	115
16.2	Installing	115
16.2.1	How do I get lpr to work properly other than for root?	115
16.2.2	Which <code>/dev/lp*</code> is my printer?	116
16.2.3	When the kernel boots it says using polling driver	116
16.2.4	What's the difference between polling and interrupt driver?	116
16.3	The Printcap Database	117
16.3.1	Where do I get a printcap for a xxxxx?	117
16.3.2	What is the content of <code>/verb"/etc/printcap"</code> ?	117
16.3.3	What is the format of a <code>/etc/printcap</code> ?	117
16.3.4	My <code>/etc/printcap</code> looks identical to someone else's	118
16.3.5	What's the minimum <code>/etc/printcap</code> ?	118
16.4	Common Problems	118
16.4.1	Why do I get errors	118
16.4.2	How do I prevent the <i>staircase effect</i> ?	118
16.4.3	Why do files with a lot of graphics get truncated?	119
16.4.4	Why doesn't <code>lpr -i</code> work?	119
16.4.5	Why doesn't <code>lpr -p</code> work?	119
16.4.6	Why does lprm say too many users ?	120
16.4.7	Why does lpc complain that it hasn't stopped a daemon?	120
16.5	Special	120
16.5.1	How do I get my printer to go back to the default font?	120
16.5.2	How do I prevent a formfeed at the end of every printout?	120

16.5.3	How do I get burst/banner pages?	121
16.5.4	How do I print text on a ps printer?	121
16.5.5	How do I print over a Network?	121
16.5.6	How do I write filters?	121
16.5.7	I've written a filter in shell script which works OK	122
16.5.8	When should I define an <code>of</code> filter?	122
16.5.9	Where do I get filters for given printers?	122
16.5.10	How do I get a program to act as a filter?	122
16.5.11	The set of filters supported by <code>lpd</code> seem strange	123
16.5.12	What are magic filters?	123
17	Linux-FAQ: UUCP/News/Mail for Linux	123
17.1	General Information	123
17.1.1	What uucp, news, and mail software works under Linux?	123
17.1.2	So where are the sources?	124
17.1.3	Do I really need the sources?	124
17.1.4	How do I get my system a news/mail/uucp feed?	124
17.2	UUCP	124
17.2.1	Why do I get 'timeout' on connections when I upgraded to uucp-1.04?	124
17.2.2	Why doesn't anonymous uucp seem to work?	125
17.3	News	125
17.3.1	What newsreader is best?	125
17.3.2	with <code>trn</code> — why do I get the message <code>dbinit failed</code> ?	125
17.3.3	Why is my <code>nn</code> database getting continuously corrupted?	125
17.3.4	Why can't I post to moderated groups?	125
17.3.5	Why do I have problems that appear to be permission-related?	125
17.3.6	Why doesn't SLS have the <i>Performance Release</i> of C-news?	125
17.3.7	Why does <code>doexpire</code> say <code>severe space shortage</code>	126
17.3.8	How do I set up NNTP?	126
17.4	Mail	126
17.4.1	Why is my <code>hostname.domainname</code> not in my headers?	126
17.4.2	Why can't <code>elm</code> write to my mailbox?	126
17.4.3	Is <code>sendmail+IDA</code> available for Linux instead of <code>smail</code> ?	126

1 Introduction

This is the introduction to a list of frequently asked questions (FAQ for short) about Linux with answers (Yeap!). This article contains a listing of the sections, followed by the question/answer part.

This FAQ is supposed to reduce the noise level ;-)) in the `comp.os.linux` newsgroup, and spare the time of many activists. I will cross-post it each month to `news.answers`. This FAQ is NOT an introduction to UNIX, there are many books for UNIX, and there is *also* a FAQ for UNIX (it's the one of `comp.unix.questions` which contains things such as "How do I remove a file named -". I DO NOT WANT TO ADD SUCH THINGS IN THIS FAQ DEVOTED TO LINUX.

The FAQ contains a LOT of information, and is pretty long (posted as many parts). Everybody should read it, 'cause it really does contain solutions to the many common problems. It can be FTP'd from at least `tsx-11` (`docs/FAQ_XXX`, where `XXX` is the date), and `nic.funet.fi` (`doc/FAQ_XXX`). Also check the archives at `pit-manager`, send mail to `mail-server@pit-manager.mit.edu` with the word `help` in it.

Some books to read:

- The C Programming Language: Kernighan & Ritchie.
- POSIX Programmer's Guide: D. Lewine
- UNIX System Administration Handbook: Nemeth, Snyder & Seebass.
- UNIX for the Impatient: Abrahams & Larson.
- UNIX System V Release 4, An Introduction, by Rosen, Rosinski and Farber; Publisher Osborne MacGraw-Hill.
- The X Windows System in a Nutshell: O'Reilly.
-

This FAQ is available at the main Linux sites in the `doc` directory, the addresses are given in section `refsecAddresses`. of this FAQ. There is also an archive of (all) FAQs at `rtfm.mit.edu` (18.172.1.27). Have a look in the anonymous ftp directory: `/pub/usenet/news.answers/linux-faq`. If you do not have anonymous ftp access, you can access the archive by mail server. Send mail to `mail-server@rtfm.mit.edu` with the words `help` and `index` in the body on separate lines for more information.

The information in this multi-parts FAQ is likely to change relatively quickly. If this is more than two months old (it was released on February 1993) then you should obtain a new copy. See the paragraph above for details of where to find a more recent version.

Please suggest any change, rephrasing, deletions, new questions, answers ...

Please include "FAQ" in the subject of messages sent to me about FAQ. Please use `linux@numero6.greco-prog.fr` whatever will be the `From` part of this message. Finally discussion about the FAQ can be done on the `DOC` Channel (see section `refsecAddresses` below).

Thanks in advance, Marc

1.1 Future Plan

- provide FAQ as diff too, since it seems to stay stable except for very few sections.
- perform automatic post to `c.o.l`, `c.o.l.a`, `c.a` and `n.a` every month as I promised long time ago.

1.2 Where to get the Linux FAQ?

PLEASE: EVERYONE read the FAQ for this group before posting questions. Many people (myself included) worked hard to produce the FAQ to accurately answer all of your questions. The new version of the FAQ (with many updates, clarifications, etc.) will only come out every month! So you won't see it posted that often. In the meantime you'll have to get it in one of the following ways:

1. ftp to `tsx-11.mit.edu`. It's the file `/pub/linux/FAQ`. It's one big ASCII file (about 150k).
2. ftp to `sunsite.unc.edu`. It's the file `/pub/Linux/docs/FAQ`. Again, one big ASCII file, same thing as on `tsx-11.mit.edu`.
3. If you don't have FTP access, you can retrieve it via the mailserver at `pit-manager.mit.edu`. Send mail to `mail-server@pit-manager.mit.edu` with the word `help` in the body.
4. If you simply can't find the FAQ, I will mail it to you myself. Just send me mail (`mdw@tc.cornell.edu`).
5. The most recent FAQ will be posted to `comp.os.linux`, `comp.os.linux.announce`, and `news.answers` every month (if you want to wait...).

2 Warnings

The FAQ contains a lot of information sometimes I've put it down in 3 different ways because people seems not to understand what they read (or what I wrote, you know I'm just a froggy and english is not my natural language). What I mean is that not all is in the FAQ but many things are there, so please just take time to read it this will spare a lot of the other linuxers [and if you think I should rephrase some Q/A just drop me a note with the corrections].

As the Linux kernel changes monthly, I define 3 pseudo variables a la C one for the version, and one for the date of the release. In the \LaTeX version of the FAQ here defs to have text in the FAQ.

```
#define CURRENT_VERSION    0.99      /* the current version */
#define PATCH_LEVEL        10        /* the patch level */
#define KERNEL_DATE        29, May   /* Date of the CURRENT_VERSION */
```

In what follows I'll consider `CURRENT_VERSION` (0.99) as the current version.

Below is the information provided by Lars Wirzenius on the LINUXNEWS Channel: "Linux News, a summary. Issue #5"

November 9 Linus announced kernel version 0.98 patchlevel 4. The most important changes are:

- The inode caching bug (resulting in bad filesystem info when mounting/umounting devices) should be gone for good.
- Bug fix of a race-condition in the filesystem (which may have caused people to get occasional fsck errors).
- Math emulator fixes (mainly for the re-entrancy problem)
- `NR_OPEN` was changed from 32 to 256 (this will break GNU Emacs, the term program, and possibly other programs; a recompilation should fix things).

- the process kernel stack is now on a separate page (needed due to
- Changes in kernel data structures (these require a new ps)
- System call tracing
- Changes to networking (tcp/ip, some nfs)
- FTP: `nic.funet.fi : pub/OS/Linux/PEOPLE/Linus/` (presumably on other sites too, by now), files `linux-0.98.4.tar.Z` (complete source), `linux-0.98.patch4.Z` (patches against pl3). Also `ps-diff.Z`, patches by Linus to `ps-0.98` that should fix `ps` to work with 0.98.4.

November 15 Linus announced kernel version 0.98 patchlevel 5.

- This version fixes a swap-partition bug in pl4. The symptoms were incorrect swapping with a partition, e.g. `xterm` could dump core when swapping was enabled and you typed at the keyboard.
- This version also checks against writing to the text segment (i.e. program code). This will break some binaries, especially some very, very old ones that were compiled using the `estdio` library (this library is no longer used in Linux). If previously well-working programs suddenly start to dump core, this might be the reason.
- There are also some other changes.
- See also 0.98.4 announcement above.
- FTP: `nic.funet.fi: pub/OS/Linux/PEOPLE/Linus/`

May 29 Linus announced kernel version 0.99 patchlevel 10.

I uploaded a new version of the ALPHA-pl10 to `nic.funet.fi`: this one contains some more fixes for the networking code, along with various other minor changes (including the math emulation fix). The directory is, as before, `pub/OS/Linux/PEOPLE/Linus`. I'd suggest removing any old traces of linux before unpacking this, so that you won't have problems with old object files etc.

Also note that Fred put the new networking binaries on `tsx-11` the other day, so you should probably get them from `/pub/linux/packages/net/net-2`. I'd suggest getting at least the `net-base` and `net-std` packages of binaries, as they contain the config programs as well as all the normal networking binaries.

I'd be especially interested to hear comments from the people who were unable to get the previous alpha-pl10 working. Does this one work for you? If it doesn't *please* mail me directly, as that problem is now the only thing keeping me from a real release.

To recap, pl10 gives you roughly:

- new net-2 code (FvK) along with the newest drivers from Donald.
- IPC doesn't need patching (Krishna Balasubramanian).
- various FPU-emulation details fixed (Bill Metzenthien)
- ext2fs updates by Remy Card and Steven Tweedie.
- dynamic inode and file allocation (Steven Tweedie) with hash tables for better inode lookup etc. Hopefully no more **EMFILE** errors.
- updated fdomain driver by Rik Faith (along with some other SCSI changes by others)

- `tty` changes by Tytso and others.
- new `uname()` call for extended info (ie domain). HLU.
- iBCS signal stacks and stubs for `lcall 7,0`.
- buffer cache / code page sharing. Small changes since the last ALPHA-diff, as people with 4MB reported that the old code didn't relinquish pages very nicely.
- updated `clone()` — it didn't actually work with `execve()` before, and there may still be some problems. Untested.
- various other changes: I've probably forgotten half the changes since 0.99pl9.

Please test it out, and if you have problems, mail me about them so that I know about it. Feel free to send them to the mailing list and newsgroup too, of course, but at least include a mail to me /- I'm working on finding the reason for the bootup problems that two persons have experienced: the more data I can get on this, the merrier.

3 Linux General Information

This section is maintained by Matt Welsh (mdw@tc.cornell.edu). Mail him if you have corrections, additions, other questions, etc. Last update was June 1993.

3.1 Information

3.1.1 What is Linux?

Linux is a small UNIX for 386-AT computers, that has the added advantage of being free ¹. It is still in beta-testing, but is slowly getting useful even for somewhat real development. The current version is `CURRENT_VERSION`, date: `KERNEL_DATE`.

Linux is a freely distributable UNIX clone. It implements a subset of System V and POSIX functionality, and contains a lot of BSD-isms. LINUX has been written from scratch, and therefore does not contain any AT&T or MINIX code—not in the kernel, the compiler, the utilities, or the libraries. For this reason it can be made available with the complete source code via anonymous FTP. LINUX runs only on 386/486 AT-bus machines; porting to non-Intel architectures is likely to be difficult, as the kernel makes extensive use of 386 memory management and task primitives.

3.1.2 Does Linux support GCC, TCP/IP, X-Windows, MGR, etc.?

Linux currently supports and uses a large amount of the GNU software (i.e. GCC, bison, groff, etc) so all of that functionality is there. X-Windows is also available, along with many client applications. MGR is there too. TCP/IP is in testing, and is available for you to try out.

Dirk Hohndel: TCP/IP is available. I use Linux boxes as Xterminals and my "own" asterix has mounted half a Gig via NFS.

Mitch DSuozza is running an anonymous FTP server on his Linux box. This is definitely more than testing. The newest SLS has a TCP/IP kernel by default.

¹Free means that you may use it, change it, redistribute it, as long as you don't change the copyright. Free does not mean public domain

See section 11!

In short, Linux supports many, many features and programs. One of the biggest questions is:

*Does ***** work on Linux? Does Linux have *****?*

The answer, usually, is "yes". Just check out the rest of this FAQ, the newsgroup, as well as the files on the FTP sites.

3.1.3 What is the current state of Linux?

read the `comp.os.linux` newsgroup, where the `Info-Sheet` is periodically posted. You can also read `comp.os.linux.announce`, which is a moderated newsgroup, has a lot less traffic, and contains all of the *important* information on Linux.

3.1.4 Linux sounds great, but where do I get it and how do I install it?

FIRST read this FAQ, and especially section 5. There is no single, *official* release of Linux — instead, there are several independent releases, all with their own advantages and disadvantages. The SLS release is the de facto standard.

Note that some releases only give you the kernel and a few utilities, and others give you everything you need (including X11, GCC, and more) in that latter case the downloading is close to a douzen of SOFT. Just check out section 5 for more info.

3.1.5 Sounds good, but all of these docs are very confusing. Is Linux really difficult to install and use?

This FAQ, as you can see, is very large. This is mostly because it's full of every single frequently asked question about Linux on the net. In essence, it's very simple: to get started with Linux, download the SLS release, put it onto floppies (see section 5), repartition your drive, and install the software.

If you're new to the UNIX world, Linux (as with any UNIX) is going to be difficult to understand at first. There isn't a lot of real documentation (other than this FAQ) out there. The Linux Doc Project (`mailmdw@tc.cornell.edu` for info) is working on a set of Linux manuals which should solve this problem. However, Linux is a hacker's UNIX, in many ways: it was developed by experienced UNIX hacks for experienced UNIX hacks. People are working on making it more *neubie-friendly*, however, keep in mind that this is an afterthought. Linux was never meant to be the huge popular free UNIX that it has become, and the lack of documentation doesn't bother UNIX wizards who can figure it out from just poking around.

3.1.6 What's the best advice you can give to a Linux newbie?

Go read a *good* book on using UNIX before you even get started. Jumping cold-turkey into the Linux world from DOS is going to be quite difficult. Also remember that you'll be running the system, as well as using it, so get a book on UNIX system administration, too.

Too many folks post questions to `comp.os.linux` which aren't Linux specific at all, and thus aren't covered in this FAQ. (i.e. *How to remove a file called -i?*). If you can't figure out something, it's more than likely a misunderstanding of UNIX concepts, not a bug in Linux. Please read up on UNIX (see the `comp.unix.questions` FAQ; that's a good one) before you get started.

3.1.7 How long has Linux been publicly available?

Few months, v0.10 went out in Nov. 91, v0.11 in Dec. and the current version 0.99 is available since 29, May. But even it is pretty recent it is quite reliable. There are very few and small bugs and in its current state it is mostly useful for people who are willing to port code and write new code. As Linux is very close to a reliable/stable system, Linus decided that v0.13 will be known as v0.95. Believe it or not: the whole story started (nearly) with two processes that printed AAAA... and BBBB... BTW consult the digest#136 Vol2 for a complete story.

3.1.8 How reliable is Linux, anyway?

Very much so. The only real *bugs* that we see are with alpha drivers (that's why they're alpha) and with some parts of the TCP/IP code. For 99% of applications, however, Linux is very robust. Linux and X-Windows is faster on a 486-33 than on many Sun workstations with the same amount of RAM, running SunOS.

3.1.9 What is the proper pronunciation for "Linux"?

Linus himself said:

'li' is pronounced with a short [ee] sound: compare prInt, mInImal etc. 'nux' is also short, non-diphthong, like in pUt. It's partly due to minix: Linux was just my working name for the thing, and as I wrote it to replace minix on my system, the result is what it is... linus' minix became Linux.

I originally intended it to be called freax (although buggix was one contender after I got fed up with some of the more persistent bugs -:) and I think the kernel makefiles up to version 0.11 had something to that effect (*Makefile for the freax kernel* in a comment). But arl called the Linux directory at `nic.funet.fi/pub/OS/Linux`, and the name stuck. Maybe just as well: freax doesn't sound too good either (freax is obviously free + freak + the obligatory -x).

(Rick's note for English speakers: Linux - "LIH-nuhks".)

3.1.10 What's about the copyright of Linux?

This is an excerpt of the RELEASE Notes v.095a: Linux is NOT public domain software, but is copyrighted by Linus Torvalds. The copyright conditions are the same as those imposed by the GNU copyleft: The GNU GENERAL PUBLIC LICENSE Version 2, June 1991 is part of the source tree.

3.1.11 Should I be a UNIX and/or a DOS wizard to install/use Linux?

Not at all, just follow the install rules, of course it will be easier for you if you know things about UNIX. Right now Linux is used by more than BIGNUM persons, very few of them enhance the kernel, some adds/ports new soft, most of us are only (but USEFUL) beta testers. Last but not least, various Linuxers work on manpages, newuser_help, file-system organization. So join us and choose your "caste".

It is even used in production environments (Dr. G.W. Wettstein)

3.1.12 Why can't we split comp.os.linux?

Ian Jackson tells us:

There is a procedure for creating new newsgroups, involving discussion periods and votes; it can be found in `news.announce.newgroups`.

In November 1992 I (Ian Jackson) started a formal discussion under that procedure and duly held a vote for four new groups, `comp.os.linux.announce` (moderated), `comp.os.linux.questions`, `comp.os.linux.bugs` and `comp.os.linux.misc`.

There was quite a heated argument, with many people (especially from Fidonet and the news→mail gateway) complaining that if the group split they wouldn't be able to read it.

At the end of the vote the results were as follows (culled from the announcement at the end of the voting period):

	yes	no	abs	diff	ratio	result why	to change
<code>.announce</code>	479	131	3	348	3.6564885	PASS	249
<code>.questions</code>	380	217	16	163	1.7511521	FAIL (ratio)	54
<code>.bugs</code>	390	212	11	178	1.8396226	FAIL (ratio)	34
<code>.misc</code>	390	207	716	183	1.8840580	FAIL (ratio)	24

diff number more yes than no votes — this must be ≥ 100 for a group to pass.

ratio ratio of yes to no votes — this must be ≥ 2 for a group to pass.

to change the minimum number of votes which would have been required to change the result (if they were all "yes" or "no" as appropriate).

The guidelines say that unless a group gets at least twice as many "yes" as "no" votes and at least 100 more "yes" than "no" votes it won't be created. Hence all the new groups except `.announce` failed. `comp.os.linux.announce` now exists: the submission address is via Matt Welsh (`linux-announce@tc.cornell.edu`).

The guidelines also say that you have to wait at least 6 months after a failed vote before trying again — this to stop the obvious problem of failed groups coming back over and over again.

Hence any more discussion of proposed splits is futile until at least very late in June. In any case, such a discussion should take place in the group reserved for that purpose, `news.groups`, not in `comp.os.linux`.

3.2 Requirements

3.2.1 Does it run on my computer?

Linux has been written on a clone-386, with IDE drives and a VGA screen. It should work on most similar setups. The harddisk should be AT-standard, and the system must be ISA. (though *some* EISA success has been reported [T. Koenig], Linux doesn't take advantage of the EISA structure). A high density floppy drive — either 5.25" or 3.5" — is required.

Drew's information: Linux supports anything that's register compatible with a WD1003 MFM disk controller (ie, the original PC-AT disk controller.) Most AT MFM, RLL, ESDI, and IDE setups look like this.

XT compatible disk controllers won't work.

Generally, the rule is if you have the disk configured into the CMOS setup of your machine, it will work (because the BIOS is talking to a WD 1003 compatible board), otherwise it won't.

IDE and MFM seem to work with no problem. It works, also, for some ESDI drive (you might have to comment out the `unexpected hd interrupt`-message from `hd.c`). There exists a high-level SCSI driver, under which low-level drivers are placed; a ST-01/ST-02 low driver has been completed see the sections 11 and 4.

Otherwise the requirements seem relatively small: a 386 (SX, DX or any 486). Any video card of the following: Hercules, CGA, EGA, (S)VGA.

It needs at least 2M to run (with SWAP), and 4M is definitely a plus. It can happily use up to 16M (and more if you want).

BTW There are problems with some MAXTOR drives on high speed machines (sometimes switching off "turbo" helps). There may also be a problem with "slow" memory (under 60ns) on fast machines. Again, the solution is to turn off "turbo". Mixed SIMMs (3 and 9 chip versions) have also reported to be problematic.

Note It doesn't run (yet?) on a MCA machine (such as the IBM PS/2 line)

3.2.2 How much space will Linux take up on my hard drive?

It depends on which release you choose. See the section 5 below. Usually it's somewhere between 10 MB (for a nominal system+swap space) and 80 MB (no longer 30-40 MB) (for everything plus space for user directories, etc.).

BTW the full SLS needs around 60 MB (including T_EX and other goodies).

3.2.3 Why the suggested 4 MB, for Linux?

Linux uses the first 640k for kernel text, kernel data and buffercache. Your mother board may eat up 384K because of the chipset. Moreover there is: init/login, a shell, update possibly other daemons. Then, while compiling there is make and gcc (2.01 770k). So you don't have enough real memory and have to page.

3.2.4 How would this operate in an OS/2 environment?

Linux will coexist with *any* other operating system(s) which respects the *standard* PC partitioning scheme - this includes DOS, OS/2, Minix etc.

WARNING: Linux and OS/2 *can* co-exist on the same machine. BUT, you cannot use Linux's `fdisk` to make Linux partitions! See the warnings in section 5 about Linux and OS/2.

3.2.5 Will Linux run on a PC or 286-AT? If not, why?

Linux uses the 386 chip protected mode functions extensively, and is a true 32-bit operating system. Thus x86 chips, x<3, will simply not run it.

3.2.6 Will Linux run on a 386 Laptop?

It works, including X on most of them.

3.3 What Linux does

3.3.1 Does Linux use TSS segments to provide multitasking?

Yes!

3.3.2 If my PC runs under Linux, is it possible to ftp, rlogin, rsh etc.. to other UNIX boxes?

Kermit and ka9q have both been ported to Linux. Also, TCP/IP is quite reliable, only a few clients are missing. Read section 14. devoted to Ethernet and Linux.

3.3.3 Does Linux do paging? Can I have virtual memory on my small machine?

Yes, it does. Generally you set up a swap file or partition, and enable it with the "swapon" command. Voila! Virtual memory.

3.3.4 Can I have tasks spanning the full 4GB of addressable 386 memory? No more 64kB limits like in coherent or standard Minix?

Since 0.97pl? it uses 4 GB Process Space, 3 for userspace and 1 for the kernel space.

3.3.5 Does the bigger program sizes mean I can run X?

Yes! See section 14 below for details on X11.

3.4 Comparing Linux with other Unices

3.4.1 What are the differences, pros and cons compared to Minix?

- Cons:
 - Linux only works on 386 and 486 processors.
 - Linux needs 2 MB of memory just to run, 4 MB to be useful.
 - Linux is a more traditional UNIX kernel, it doesn't use message passing.
- Pros:
 - Linux is free, and freely distributable, BUT copyrighted.
 - Linux has some advanced features such as:
 - * Memory paging with copy-on-write
 - * Demand loading of executables
 - * Page sharing of executables
 - * Multi-threaded file system
 - * job control and virtual memory, virtual consoles and pseudo-ttys.
 - Linux is a more traditional UNIX kernel, it doesn't use message passing.

3.4.2 What are the pros and cons compared to 386BSD ?

Linux and 386BSD started out as completely different projects, with completely different goals and design criteria in mind. Personally, my experience with 386BSD has shown that :

- its TCP/IP is more stable, and seems to work (on my NE2000 card, which Linux doesn't support yet), but
- the interface and *feel* just weren't the same as Linux (which I liked better :). You'll have to try it out for yourself. Here's information from `comp.unix.bsd`, September 1992:

This is a free BSD for the 386. It does NFS and X. Grab it from `agate.berkeley.edu` (128.32.136.1). Bill Jolitz is writing a book on it.

It has a neat, small kernel, and is stable enough to run for a week or more without crashing/hanging. It works fine with 4 MB RAM and 80 MB disk. Running X however needs 8MB RAM, and supports only vanilla VGA. The TCP/IP implementation in this is functional, but not optimised. This means that FTP throughput is about 5KB/sec on Ethernet, when almost anything else gives you about 70KB to 100KB per second.

Read `comp.unix.bsd`. That is almost entirely Jolitz's 386BSD now.

1. How much disk space for:

- kernel
- X
- source to kernel
- source to X
 - Binary only: 49 M (30 disk, 5 swap, 14 extraction)
 - Source: 125 M (49 bin, 37 disk, +6 extraction, 25 recompile)
 - User software: +51 M

X Binaries require about 61 MB (?) I don't know about the sources.

To conclude this Q/A wrt Linux vs 386BSD, you should read the Linux News Issue #5 of L. Wirzenius, where a true story of Mark W. Eichin is reported.

There are newsgroups devoted to 386BSD : `comp.os.386bsd.*`

- I have seen in may the first attempt for 386BSD FAQ. Nevertheless the Linux FAQ is not bug free, and contains some outdated information.
- 386BSD can do POSIX and BSD
- Linux can do POSIX, SYSV and some BSD stuff
- Linux was developed with portability in mind. So it's not directly System V, nor is it directly BSD. It's the best of both worlds, a la SunOS. :)

For most of the *nix* users both systems are fairly usable, but none of them are bug free.

4 Linux useful Addresses

4.1 Linux on the Net

4.1.1 Where can I get Linux?

Linux (all the software, binaries, sources, releases, and so on), can be retrieved via anonymous FTP from :

- EUROPE
 - `nic.funet.fi` (128.214.6.100)
directory `/pub/OS/Linux`
 - `ftp.informatik.tu-muenchen.de` (131.159.0.110)
directory `/pub/Linux`
 - `src.doc.ic.ac.uk` (146.169.3.7)
directory `/pub/os/Linux`
 - `ftp.mcc.ac.uk` (130.88.200.7)
directory `/pub/linux`
 - `ftp.dfv.rwth-aachen.de` (137.226.4.105)
directory `/pub/linux`
 - `ftp.informatik.rwth-aachen.de` (137.226.112.172)
directory `/pub/Linux`
 - `ftp.ibr.cs.tu-bs.de` (134.169.34.15)
directory `/pub/os/linux`

- US
 - `tsx-11.mit.edu` (18.172.1.2)
directory `/pub/linux`
 - `sunsite.unc.edu` (152.2.22.81)
directory `/pub/Linux`
 - `wustl.wuarchive.edu` (128.252.135.4)
directory `/pub/mirrors4/linux`
 - `ftp.eecs.umich.edu` (141.212.99.7)
directory `linux`
- AUSTRALIA
 - `kirk.bu.oz.au` (131.244.1.1)
directory `/pub/OS/Linux`
- JAPAN
 - `utsun.s.u-tokyo.ac.jp` (133.11.11.11)
directory `/misc/linux`
 - `kuis.kyoto-u.ac.jp` (130.54.20.1)
directory `/Linux` and `/Linux/mirror` for the tsx mirror
- KOREA
 - `cair.kaist.ac.kr` (143.248.11.170)
directory `/pub/Linux`, a mirror of sunsite

You might want to check out which of these is the most up-to-date. Note that `banjo.concert.net` is no longer a Linux FTP site :(

Lee M J McLoughlin (`lmjm@doc.ic.ac.uk`) said;

We (`src.doc.ic.ac.uk`) are also on Janet (the main UK academic network) as `uk.ac.ic.doc.src` (000005102000).

More useful perhaps is we are the only big archive available via FTAM, the ISO equivalent to FTP. We can be reached either over the internet or janet (see above addresses) or via the European IXI network on 204334504108

If you have no FTP capability, you are in trouble. See the next Q/A. Also, you'll need the `UNCOMP.EXE` and `RAWRITE2.EXE` programs for DOS (to make your install disks). These are usually found in the Linux directories on the above FTP sites.

4.1.2 I do not have FTP access, what can I do to get Linux?

You can either read the next subsection related to BBS's otherwise, read the following.

The SLS release is distributable by snail-mail on floppies for those without net access; see section 5.2.3 of this FAQ for more.

Try to contact a friend on the net with those access, or try mailserver/ftpmail server otherwise contact `tytso@ATHENA.MIT.EDU`. You might try mailing `mailserver@nic.funet.fi` with "help" in the body of the mail. If you choose ftpmail server (example: `ftpmail@decwrl.dec.com`), with "help" in the body, the server will send back instructions and command list. As an example to get the list of files available at `tsx-11` in `/pub/linux` send:

```
mail ftpmail@decwrl.dec.com
subject: anything
reply $<your e-mail>$
connect tsx-11.mit.edu
chdir /pub/linux
dir -R
quit
```

In Europe `ftp.informatik.tu-muenchen.de` is accessible via e-mail (send `help` in the body to `ftp-mailer@informatik.tu-muenchen.de`)

4.1.3 What is a quick way to find out who is offering mail-order Linux kits?

If you have access to e-mail, e-mail `bcr@physics.purdue.edu` a message with the subject: `help`.

4.1.4 Is there a newsgroup or mailing-list about Linux? Where can I get my questions answered? How about bug-reports?

The `comp.os.linux` newsgroup is literally *teeming* with postings. So, to the first question, yes. :) The older newsgroup, `alt.os.linux`, is being phased out and shouldn't be used anymore.

There are other newsgroups, because not all problems you encounter while using Linux are Linux specific. You may also want to read several groups in the `comp.unix.hierarchy`, `comp.sources.wanted`, `alt.sources.wanted`, and especially any FAQs in these groups.

If you don't have news access you can get the digest of postings via e-mail from: `linux-activists-request@news-digests.mit.edu`. This list is gatewayed to the newsgroup as well. Only use the 'request' address for subscribe/unsubscribe messages; don't post those to the newsgroup or to the actual mailing list.

Questions and bug-reports can be sent either to the newsgroup or to Linus himself at `torvalds@kruuna.helsinki.fi`, depending on which you find more appropriate.

And last but not least there is the original mailing-list, which is now a multi-channel list. Contact `linux-activists-request@niksula.hut.fi`.

4.1.5 Where can I get my questions answered? How about bug-reports? What do I put into a post to `comp.os.linux`?

Paul Gortmaker (`pg@cain.mmtc.rmit.oz.au`) says about it:

You can post your problem to the above group, `comp.os.linux`. BUT, BEFORE YOU DO THIS, PLEASE READ THE FOLLOWING GUIDELINES.

If you have read the FAQ, man pages, etc, and you still haven't solved your problem, then check to make sure you have got the latest version of whatever it is that you are working with. Check the dates and revision numbers of your versions with the versions on your local ftp site (`tsx-11.mit.edu`, `sunsite.unc.edu`?). This includes (most importantly) the kernel itself. Make sure you have applied the latest patches and recompiled the kernel, or have got the kernel **Image** from someone who has done so. And, of course check `comp.os.linux` for info too. If you have a genuine problem, chances are that you aren't the first one to find it. So it has probably already been reported (...and fixed???). For example, if you are having trouble with say Xconfig for some strange VGA card, and you use `nn` to read `comp.os.linux`, then you could invoke `nn` as follows:


```
nn -x -s 'config' comp.os.linux
```

and it will find all the latest articles with the word config in their subject for you. This will be one of the most up to date sets of information that you can get – DON'T OVERLOOK IT !!! (You can check the man pages of your news reader to determine the options that do the same as the above.)

OK, so you've done all the above, spent 40 hours trying to figure it out, have had a nervous breakdown, your girlfriend/boyfriend has stopped talking to you, and you decide that you will turn to the Linux community for help. Here are some guidelines on posting that will ensure that you get a quick response, and that you hopefully don't get flamed.

1. Choosing a Subject:

It is important to try and squeeze as much information into as few words as possible. If you can manage it, try and put the package name, version, and problem into the subject. But don't make it too long, or the middle will get chopped out. For example "I'm having problems with poeig-1.1.tar.Z on my 486 with 0.99p6" will probably appear to everybody as "I'm having prob with 0.99p6" Not very useful... What should have been used was something like: "poeig-1.1 w 99p6 wont compile" would be much better, and relays that you are having trouble with getting it to compile. (Note that this is just an example, I have no knowledge of problems with poeig!) Also, (unless you like bugs → getting flamed!) DON'T claim you have found a bug, unless you are ABSOLUTELY SURE! Nothing p***es developers off more than erroneous bug reports.

2. Keywords:

If your news poster program asks for keywords, try and put in some useful descriptive words, so that others can use them for a meaningful search.

3. Body of the Article:

There are some key things that need to be included in the body of the article.

- The name and version of the thing that you are having the problem with.
- The type of problem, ie compilation, execution, etc.
- versions of related software, ie if compilation is the problem, then the version of GCC you are running is relevant. If you are having trouble with a program that uses X, then the version of X you are using is relevant.
- The version and patchlevel of the kernel you are using at present. (ie. 0.99p7 or whatever)
- the type or brand of any related hardware, ie. if you are having problems with networking, then you would want to say that you are using a Western Digital SMC Elite 16 or whatever your ethernet card is.
- Any relevant error messages that were reported by the system during the problem.

And, of course, there are some things that one should **not** put in the article. For example, don't post a 30 page configuration file and expect anybody in their right mind to look through it. And similarly for any HUGE files. If they are relevant to your problem, then someone will respond by asking you something like "Did you check line 32 in file such and such???" And try to

avoid negative comments like "The documentation isn't fit for my dog." If you have a bone to pick, do it via e-mail, so the rest of us don't have to read a flame war! It just adds to the amount of useless noise on `comp.os.linux`, which already takes too long to scan through. Besides, the developers are doing this FOR FREE. THEY ARE NOT OBLIGATED TO DO ANYTHING. DON'T ABUSE THEM!!! (Or they might just go away, which hurts us all.)

Well, with all this in mind, hopefully you will get a quick response to your problem, and maybe someday you will be able to answer someone else's problem from the experience you gain!

4.1.6 Could you be more explicit about the multi-channel list?

Well, there are many things to say:

- these channels are rather devoted to hackers
- the ones I am aware of are: GCC, MGR, X11, SCSI, NEW-CHANNELS, MSDOS (emulator discussion) , NORMAL, KERNEL, FTP, LAPTOP, DOC, NET, CONFIGS, LINUXNEWS ...
- whenever you want to JOIN or LEAVE a channel you have to contact the request address
- you have to use special header (**X-Mn-Key** and **X-Mn-Admin**); **X-Mn-Key** is *only* for regular post, the **X-Mn-Admin** is for *request*

(Ari Lemmke: 1 Nov. 1992): Hmmmm.... It seems our list has now about 1500 users in 21 channels (mailing lists). 3960 without uniq.

channel info	231
user info	150
failed to join	48
failed to leave	126
unknown command	246
no Mail-Net field	351
left the channel	388
notice sent	687
articles to digests	877
digests sent	659
users on all channels	6739
uniq users on	2725

```
echo foo | mail linux-activists-request@niksula.hut.fi
```

to get the Mail-Net User Guide.

4.1.7 How can I join the channel XXX on the linux-activists mailing list?

just send a mail to the request address with help in the body; you will get back a mail which gives you the list of channels and the way to join/leave them. Basically you send mail to the request address with the line:

```
X-Mn-Admin: join <channel>
```

4.1.8 How can I leave the channel XXX on the linux-activists mailing list?

Same as above, basically. You send mail to the request address that contains the line:

```
X-Mn-Admin: leave <channel>
```

4.1.9 I'm not an hacker, what are the channels I could be interested in?

Probably these are the most interesting for you (IMHO)

Channel NORMAL: Channel normal is the former Linux-Activists mailing list (all the people who were on the old Linux-Activists list are moved to this channel).

Channel DOC: This channel is for Linux document "project". Discussion about Linux documents, manuals, papers, etc.

Channel CONFIGS: This channel will be devoted to send submissions of systems that have Linux already running, AND those that, for any reason, can't get it to work yet.

Mainly, what it's need from all the channel users is to send their hardware configuration list (as complete as possible). Include anything that you feel pertinent for information: CPU, motherboard, RAM amount, HD & floppy controller, BIOS, monitor, video card & memory, network adapter, etc. If you are having trouble with your current system, or you find out that a program doesn't work properly on your system due to a HARDWARE problem, it may be useful for us to know your configuration; maybe you can get a lot of help.

Channel LINUXNEWS: The LINUXNEWS channel will be used for distribution of Linux News, a weekly (if I can find the time) summary of things that happen in the Linux community. Discussion is not encouraged, if you have complaints or suggestions, send them directly to me (Lars.Wirzenius@Helsinki.FI).

Channel ANNOUNCE: This channel contains the hole traffic of `comp.os.linux.announce`. Every article you send to this group will be sent out as mail in the ANNOUNCE channel. This is very usefull for those who doesn't have news access and those who can't afford to spend one to three hours per day reading news, but want to be up to date.

Channel NEW-CHANNELS: On the future users on this channel get the information about new channels created.

By this way you can join the channels you want, and do not need to send mail to Mail-Net info server or listen rumours.

4.1.10 Does there exist a place where the traffic of the newsgroup is kept?

Yes, on nic and tsx-11 (see the ftp addresses above), and since 12th March, a Gopher server is up at `bery1.daimi.aau.dk` (130.225.16.86). The archives go back to Nov. 18. 91. Also recently a WAIS server for the Linux mail archive has been setup at `fgb1.fgb.mw.tu-muenchen.de`. Contact tw@fgb1.fgb.mw.tu-muenchen.de for more info.

All back issues of the Digest are available on `tsx-11.mit.edu` (18.172.1.2) in the following place(s):

`/pub/linux/mail-archive/Volume[x]/digest[nnn].Z`

where `[x]` is volume `x` and `[nnn]` is issue # `nnn`

4.2 Obtaining Linux from BBS's

This section is maintained by Zane Healy (`healyzh@holonet.net`)

Last Update June 1993.

4.2.1 I don't have access to FTP, how can I obtain Linux?

Linux is available from various BBS's around the world.

4.2.2 I got this FAQ from a local BBS, or a friend, and I see there is a newsgroup called `comp.os.linux`. I don't have access to USENET or mail, so how can I get the messages?

Some of the BBS's on FidoNet carry `comp.os.linux` as a FidoNet conference. Also some of the other BBS's carry it in some form or other.

4.2.3 Do BBS's offer anything that the Internet does not?

Yes, on the information side there are the UNIX conferences on both the RIME network and FidoNet. Although they are not dedicated to Linux, a large amount of the messages are Linux related. Also at least one software package being developed for Linux, and also one port is available via BBS's long before they are available via anonymous FTP.

4.2.4 What is a BBS?

A BBS is a Bulletin Board System, it let's you transfer message's and file's via your phone line and all you need is a computer with communications software and a modem. Some BBS's transfer message's among each other forming large computer network's similar to USENET. The most popular of these in the US are FidoNet and RIME.

4.2.5 How can I get a (Near) complete list of BBS's that carry Linux?

I (Zane Healy) post a list of all known BBS's that carry Linux to `comp.os.linux` as well as the RIME and Fidonet UNIX conferences on the 1st and 15th of each month.

4.2.6 Now that I have a phone number, how do I go about accessing a BBS?

1. You need a computer equipped with communications software and a modem.
2. For ALMOST all BBS's you will need to set the comm software up for:
 - 8 - Data Bits
 - N - Parity
 - 1 - Stop Bit
3. Although certain BBS's and Communication services require that the software be set for:

- 7 - Data Bits
- E - Parity
- 1 - Stop Bit

You will also need to set the comm software for the correct speed, either the top speed, or the max speed for your modem.

4. Using the comm software, call the BBS. Once you connect with the BBS (this may take awhile, as other people are likely to be using it), you will be asked some questions.

If you are a registered user of the BBS it will normally only ask for your name and password. However if you are not a registered user, it will most likely require that before you do anything, you register. The method of registration varies from BBS to BBS.

Normally the first thing that will happen is, you sign on to the BBS, and tell it your name. It will then check it's list of user's and see that you are not one of them. At which time it will ask you if you are a new user, or if you wish to re-enter your name. When you tell it you are a new user, it will then ask you some questions about yourself, such as where you are calling from and your phone number. A lot of BBS's will want some statistical info such as what type of computer you are using, your communications software, your age, etc, etc.

After this, most BBS's require some sort of validation, this is for the System Operator's (SysOp's) protection. One type is where you give the BBS software your phone number, hang up, and the BBS calls your computer to verify that you gave it legitament phone number. Some BBS's require that you mail the SysOp a postcard. Most, however just require that you give the SysOp the request info and then he upgrades your level of access a couple of day's later.

For the most part you will find that the registration process is easy to follow and well documented.

4.2.7 There is a local BBS that carries Linux, but it isn't on the latest Linux BBS List. How do I go about submitting it for inclusion in the list?

Send the following information on the BBS to me:

- BBS Name
- Phone Number
- Modem Speed
- City and State/Country
- Whatever Network it's on (i.e. FidoNet, RIME, etc.)
- First Time access to D/L Linux Files (Y/N)
- Free Access to Linux Files (Y/N)
- Allow File Requests (Y/N)
- BBS Rating (1-5)

I can be reached at one of the following E-Mail Addresses:

Internet: healyzh@holonet.net

CompuServe: 70332,14

Prodigy: SCNN49A

Fido NetMail: Zane Healy at 1:109/615

RIME UNIX Conference: Zane Healy

4.2.8 What can I do to help ensure the continued development of Linux?

PLEASE UPLOAD FILES TO BBS'S

I would like to point out that a very large number of the Linux enthusiasts don't have FTP access. In fact it is possible that by now most of the Linux fan's don't. So I would like to suggest that those of us that do, find at least one BBS to post the Linux file's to. I, for one post every file that I get to at least one of the local BBS's, and from there they the file's tend to find there way to other local BBS's. I've seen posts about the future of Linux etc., well here is a way to help guarantee it. I think it's safe to assume that most people with FTP access also have a modem. So how about doing other Linux fan's a favor and finding a BBS to upload the Linux files to.

4.2.9 How do I read the data contained in the Linux BBS List?

The list uses the following format:

State	YYY	BBS Name	Phone Number	Modem Speed
Rating		City	Other data	

The BBS's are rated by the number of Linux related file's that they carry. This is so you can choose one's that has a better chance of carrying the file's you are looking for if you are calling long distance. The BBS's are rated on a scale of one to five.

1. Only enough the most basic of files
2. The basics and a little more
3. So, so
4. A respectable amount
5. Pretty much everything you need

Information about the boards access policies can be obtained by checking a three digit field.

```

YYY -- Either a Yes/No/? answer to the question
|||
||Free access to Linux files
|Allow file requests (FidoNet)
First time D/L of Linux related files

```

NOTE: Just because a board has **N**'s in the first two fields does not mean that it is a board to stay away from. A lot of boards require that you register and be verified before you can access most of their features, hence the first **N**. The second field is, to the best of my knowledge, limited to BBS's that are part of FidoNet.

4.2.10 What are some of the best BBS's to check out?

In the US:

```
CA  YNY hip-hop          408-773-0768 14.4k V.32bis/HST
   5  Sunnyvale          Login: guest (no password)
```

For more Information look at the BBS-List.

BBS OF THE MONTH:

```
MD      Brodmann's Place      301-843-5732   14.4k
  This BBS not only carries an excellent selection of Linux file, but
it also has the following Linux related newsgroups/conferences.
  USENET          Fidonet          RIME
  -----          -
comp.os.linux      Fido UNIX          RIME UNIX
comp.os.linux.announce Linux_Local(local to 109 Fido area)
```

FORMAT:

```
State  YY  BBS Name          Phone Number   Modem Speed
Rating   City          Other data
```

```
RATING SYSTEM:  1 -- Only enough the most basic of files
                 2 -- The basic's and a little more
                 3 -- So, so
                 4 -- A respectable amount
                 5 -- Pretty much everything you need
```

ADDITIONAL INFO:

```
YYY -- Either a Yes/No/? answer to the question
|||
||Free access to Linux files
|Allow file requests (FidoNet)
First time D/L of Linux related files
```

File Requests:

FidoNet BBS's with the right type's of front-end mailer's can call other Fido BBS's and request their front-end mailer to send them files that they want. All this can be done automatically. File Requests (freqs) are basically the FidoNet equivalent to UUCP.

UNITED STATES:

```
??  ??? ??????????????????????  516-244-7064   9600
   3-4  ??????  ????
```

```
CA  NNY Citrus Grove Public Access 916-381-5822 ZyXEL 16.8/14.4
   3  Sacramento          citrus.sac.ca.us
CA  High[er] Powered BBS          408-737-7040
   4  ?                    RIME ->HIGHER
CA  hip-hop          408-773-0768   19.2k
          408-737-8300   38.4k
```

5	Sunnyvale	USENET	
CA	YYY Unix Online	707-765-4631	9600
4	Petaluma	USENET	
CA	The Outer Rim	805-252-6342	
?	Santa Clarita		
CA	Programmer's Exchange	818-444-3507	
		818-579-9711	
?	El Monte	FidoNet	
CA	??? Micro Oasis	510-895-5985	14.4k
?	San Leandro		
CA	YNY Test Engineering	916-928-0504	
?	Sacramento		
FL	NYN Slut Club	813-975-2603	USR/DS 16.8K HST/14.4K
5	Tampa	FidoNet 1:377/42	
FL	Lost City Atlantis	904-727-9334	14.4k
4	Jacksonville	FidoNet	
FL	YYY Acquired Knowledge BBS	305-720-3669	14.4k v.32bis
5	Fort Lauderdale	Internet (UUCP)	
GA	AVSync	404-320-6202	
2	Atlanta		
GA	YYY Information Overload	404-471-1549	19.2k ZyXEL
5		FidoNet 1:133/308	
GA	Atlanta Radio Club	404-850-0546	9600
?	Atlanta		
ID	Rebel BBS	208-887-3937	9600
5	Boise		
ID	YYY Phantasia BBS	208-939-1350	9600
5	Boise	Smartnet 1:347/25	
ID	YYY Rocky Mountain HUB BBS	208-232-3405	38.4k
4	Pocatello	Fido,SLNet,CinemaNet,etc	
IL	YYY EchoMania BBS	618-233-1659	14.4k HST
3	Belleville	Fido 1:2250/1 (f'reg LINUX)	
	F'reqs from unlisted nodes, online callback verifire (works L.D.)		
IL	YNY UNIX USER	708-879-8633	14.4k
4	Batavia	USENET, Internet mail	
	Home of Unix User newsletter		
IL	NYN PBS BBS	309-662-2042 Node 1 - 16.8k v.32/HST	
		309-663-7675 Node 2 - 2400	
2	Bloomington	Fido 1:232/303	
IL	Third World	217-356-9512	9600 v.32
3-4			
IN	NYN Digital Underground	812-941-9427	14.4k v.32bis
5	USENET News Feed		
LA	The OA Southern Star	504-885-5928	
?	New Orleans	FidoNet 1:396/1	
MA	N?N Channel One	617-354-8873	
?	Boston	RIME ->CHANNEL	
MA	YNY VWIS Linux Support BBS	508-793-9568	9600
4	Worcester		
MA	YYY WayStar	(508)481-7293	14.4k v.32bis
	(508)481-7147	(508)480-8371	
5	Marlborough	FidoNet (1:322/140)	
MD	N?N Programmer's Corner	301-596-1180	9600
5	Columbia	RIME	


```

MD      Brodmann's Place      301-843-5732    14.4k
  5      Waldorf              RIME ->BRODMANN
                          FidoNet
    Carries 5 different message groups dealing with Linux/UNIX
MD      Main Frame           301-654-2554    9600
  4      Gaithersburg        RIME ->MAINFRAM
MD      1 Zero Cybernet BBS   301-589-4064
  2
MD      YNY WaterDeep BBS     410-614-2190    9600 v.32
  5      Baltimore
ME      ??? Harbor Heights BBS 207-663-0391
  ?      Boothbay Harbor
MN      YNY Part-Time BBS     612-544-5552    14.4k v.32bis
  ?      Plymouth
MO      NNY The Sole Survivor 314-845-6616    14.4k v.32bis
  5      St. Louis           WWIVnet, WWIVlink, +more
NC      MAC's Place           919-891-1111    16.8k, DS modem
  5      Dunn                RIME ->MAC
NC      YNY Digital Designs   919-423-4216    14.4k,23k
  4      Hope Mills
NE      Flite Line           402-421-2434
  2      Lincoln            RIME ->FLITE    DS modem
NE      Legend               402-438-2433
  2      Lincoln            DS modem
NE      MegaByte Mansion     402-551-8681    14.4 V,32bis
  ?      Omaha
NJ      Mycroft QNX          201-858-3429    14.4k
  4      ?
NJ      YNY Steve Leon's     201-886-8041    14.4k
  3      Cliffside Park
NJ      YYY Dwight-Englewood BBS 201-569-3543    9600 v.42
  3      Englewood, NJ      USENET
NJ      YNY WEFUNK, The Mothership Connection 908-940-1012    38.4k
  4      Franklin Park, NJ
NY      YYY The Laboratory   212-927-4980    16.8k HST, 14.4k v.32bis
  3-4      FidoNet 1:278/707
OR      YYY Intermittent Connection 503-344-9838    14.4k HST v.32bis
  5      Eugene, Ore        1:152/35
    f'req LINUX for a list - CBV not needed to d/l linux files
PA      NNY Centre Programmers Unit 814-353-0566    14.4k V.32bis/HST
  5      Bellefonte, PA
PA      YNY Allentown Technical 215-432-5699    9600 v.32/v.42bis
  4      Allentown          WWIVNet 2578
PA      YYY Tactical-Operations 814-861-7637    14.4k V32bis/V42bis
  1      State College      Fidonet 1:129/226 tac_ops.UUCP
TX      North Shore BBS     713-251-9757
  2      Houston
TX      The Annex           512-575-1188    9600 HST
  ?      Fidonet 1:3802/217
          512-575-0667    2400
          Fidonet 1:3802/216
    Files available by f'req or by sign-on with a 60-minute per-day limit
TX      Walt Fairs          713-947-9866
  2      Houston           FidoNet 1:106/18

```

TX	YYY CyberVille	817-249-6261	9600
	3	FidoNet 1:130/78	
TX	YNY splat-ooH	512-578-2720	14.4k
		512-578-5436	
	5	Victoria	
TX	YNY alaree	512-575-5554	14.4k
	5	Victoria	
TX	YNY Ronin BBS	214-938-2840	14.4 HST/DS
	2	Waxahachie (Dallas)	RIME,Intelec,Smartnet,and more!
VA	VTBBS	703-231-7498	
	5	Blacksburg	
VA	MBT	703-953-0640	
	?	Blacksburg	
VA	NOVA	703-323-3321	9600
	4	Annandale	Fidonet 1:109/305
VA	Rem-Jem	703-503-9410	9600
	2	Fairfax	
VA	Enlightend	703-370-9528	14.4k
	3	Alexandria	Fidonet 1:109/615
VA	YYY My UnKnown BBS	703-780-6890	14.4k V.32bis
	5	Fidonet 1:109/370	
VA	YN? Georgia Peach BBS	804-727-0399	14.4k
	1	Newport News	
WA	YYY S'Qually Holler	206-235-0270	14.4k USR D/S
	5	Renton	FidoNet: 1:343/34
		USENET squally.halcyon.com	
		(Anon UUCP: nuucp nuucp /sc2/bbs/filelist.z)	
WA	YYY Top Hat BBS	206-244-9661	14.4k
	2	Fidonet 1:343/40	
WA	YNY victrola.sea.wa.us	206-838-7456	19.2k
	3	Federal Way	USENET

Commercial Services that carry some of Linux:

--	NNN Compuserve CIS	????????????	9600
	1	UnixForum	800-848-8199 voice number
--	NNN GENie	????????????	9600/2400
	3	800-638-9636 voice number	
		301-251-6415 voice number,	international
--	NNN BIX	????????????	
	1	800-227-2983 voice number	
--	NNN Delphi	????????????	
	?	In PC SIG	

OUTSIDE US:

AUSTRIA:

W	YYY Galaktische Archive	0043-222-8303804	16.8 ZYX (19:00-7:00)
	4	Wien	fido 2:310/77

AUSTRALIA:

NSW	YYN Linux-Support-Oz	+61-2-418-8750	v.32bis 14.4k
	5	Sydney	Internet/Usenet, E-Mail/News
NSW	YYN 500cc Formula 1 BBS	+61-2-550-4317	V.32bis

4 Sydney

CANADA:

AB	NNN	Magic BBS	403-569-2882	14.4k	HST/Telebit/MNP
	3	Calgary, AB, Canada			Internet/Usenet
AB	Y?Y	Logical Solutions	2400 Baud lines -	299-9900 to 9911	
	5		14.4 K lines	- 299-9912 to 9913	
			16.8k USR v32bis-	299-9914 to 9917	
AB	YNY	V.A.L.I.S.	403-478-1281	14.4k	v.32bis
	5	Edmonton			USENET
ON	???	The Windsor Download	(519)-973-9330	v32bis	14.4
	?				
ON	YYY	r-node	416-249-5366	2400	
	3	Toronto			USENET
QC		Synapse	819-246-2344	819-561-5268	
	4	Gatineau			RIME->SYNAPSE
QC	YNY	Radio Free Nyongwa	514-284-6693	v.32bis	(ZyXEL)
	2	Montreal			USENET, Fido

GERMANY:

The BBS's in the German and Austrian sections are thanks to Rasca Gmelch who maintains a list of German and Austrian BBS's. He can be reached at: rasca@marie.physik.tu-berlin.de

HB		bakunin.north.de	(0421) 870532	9600	
	?	D 2800 Bremen	kraehe@bakunin.north.de		
NDS	???	DataComm1	0531/13216	14.4	HST
	?	Braunschweig			fido 2:240/550, LinuxNet
NDS	???	DataComm2	0531/13217	14.4	HST
	?	Braunschweig			fido 2:240/551, LinuxNet
NDS	YYY	Linux Server /Braukmann	0441/592963	16.8	ZYX
	5	Oldenburg			fido 2:241/2012, LinuxNet
HH	???	Hub Hamburg & trash-hh	040/?	14.4	(8:00 - 24:00)
	?	Hamburg			fido 2:241/2100, LinuxNet
NDS	YYY	MM's Spielebox	05323/3515	14.4	ZYX
	5	Clausthal-Zfd.			fido 2:241/3420, SLS 1.01, SLT
NDS	YYY	MM's Spielebox	05323/3516	16.8	ZYX
	5	Clausthal-Zfd.			fido 2:241/3421, SLS 1.01, SLT
NDS	YYY	MM's Spielebox	05323/3540	9.6	
	5	Clausthal-Zfd.			fido 2:241/3422, SLS 1.01, SLT
NDS	YYY	Bit-Company	05323/2539	16.8	ZYX MO
	5	Clausthal-Zfd.			fido 2:241/3430, SLS 1.01
BW	NYN	Fractal Zone BBS /Maass	0721/863066	16.8	ZYX (24h)
	3	Karlsruhe			fido 2:241/7462
NRW	???	Hipposoft /M. Junius	0241/875090	14.4	HST (4:30-7,8-23:30)
	5	Aachen			fido 2:242/6, SLS1.01/kernel 0.99.9
???	YYY	UB-HOFF /A. Hoffmann	0203/584155	19.2	ZYX+
	3	Duisburg			fido 2:242/37, SLS1.0/kernel 0.99.7
SHL	???	FORMEL-Box	04191/2846	16.8	ZYX (6:00-20:00)
	?	Kaltenkirchen			fido 2:242/329, LinuxNet
???	???	(boxname)	?	16.8	ZYX (10:00 - 22:00)
	?	?			fido 2:246/55.4

BAY ??? (boxname) 08161/82615 16.8 ZYX (22:00-8:00)
 ? Freising fido 2:246/129
 BAY ??? BOX/2 089/6019677 16.8 ZYX (22-24,0:30-2,5-8)
 ? Muenchen fido 2:246/147, info magic: LINUX
 BAY YYY DBP Line 2+1 0851/55596 14.4 V32b (8:00-3:30)
 2 Passau fido 2:246/200
 BAY YYY DBP Line 1 0851/753789 16.8 ZYX (8:00-3:30)
 2 Passau fido 2:246/2000
 BAY YYY DBP Line 3 0851/73273 14.4 HST (5:00-3:30)
 2 Passau fido 2:246/202
 BAY YYY DBP ISDN 0851/950464 38.4/64k (V.110/X.75)
 2 Passau fido 2:246/201 (8:00-24:00,1:00-3:30)
 BLM ??? (boxname) 030/6866250 16.8 ZYX
 ? ? fido 2:2403/17
 BLM YYY CS-Port 030/4913418 19.2 ZYX+
 4 Berlin fido 2:2403/13, SLS1.02
 BLM YYY BigBrother / R. Gmelch 030/3356328 16.8 Z16 (16:00-23:00)
 5 Berlin fido 2:2403/36.4, SLS1.02/kernel 0.99.9
 BW YYY Echoblaster BBS #1 07142/21392 HST/V32b (7-19:00,23-01h)
 5 Bietigheim fido 2:2407/4, LinuxNet
 BW YYY Echoblaster BBS #2 07142/21235 V32b (20:00-6:00)
 5 Bietigheim fido 2:2407/40, LinuxNet
 BW NYN LinuxServer 0711/756275 16.8 HST (8:3-17:5,19-2) MO
 5 Stuttgart fido 2:2407/34, LinuxNet
 BW NYY Rising Sun BBS 07147/3845 16.8 ZYX (05:30-02:30)
 4 Sachsenheim fido 2:2407/41, LinuxNet

FINLAND:

NNY The Field of Inverse Chaos +358 0 506 1836 14.4k v32bis/HST
 4 Helsinki, Finland USENET; ichaos.nullnet.fi

FRANCE:

NNY Modula BBS +33-1 4043 0124, +33-1 4530 1248 HST 14.4 V.32bis
 5 Paris Michel Parlebas (no fee for Linux files)
 NNY Windows Manor ?????????????? NEED MORE INFO PLEASE.
 ? Paris Francis Rozange (recently relocated)
 YYY BuBullux (semi-private, number by request) 16.8k V32bis
 5 Paris send requests for number to perrier@onera.fr
 free semi-private system, hours 18:30-08:00 (but weekends 24 hours)
 2 hours download time per day, FREQ for both listed and
 unlisted nodes.
 NYY STDIN BBS +33-72375139 V32bis
 5 Lyon, Laurent Cas FidoNet 2:323/8
 NYY Le Lien +33-72089879 HST 14.4/V32bis
 ? Lyon, Pascal Valette FidoNet 2:323/5
 YNY Basil +33-1-44670844 V32bis
 2 Paris, Laurent Chemla
 BBS under Linux (xbbs)
 YNY Cafard Naum +33-51701632 V32bis
 2 Nantes, Yann Dupont
 open between 08:30-22:30 local time - BBS under Linux (pbbs 1.9)

IRELAND:

NYN TOPPSI +353-1-711047 9600 + HST

		+353-1-773547	14.4k v.32bis
4	Dublin, Ireland	Fido, (Chatnet ?)	Fidonet 2:263/151
	NNN DUBBS	+353-1-6789000	19.2 ZyXEL
2	Dublin, Ireland	Fidonet 2:263/167	
	NNN Galway Online	+353-91-27454	14.4k v32b
4	Galway, Ireland	RIME, @iol.ie	
	N?Y Nemesis' Dungeon	+353-1-324755 or 326900	14.4k v32bis
4	Dublin	Fidonet 2:263/150	
ITALY:			
	NYN nonsolosoftware	+39 51 6140772 v.32bis,v.42bis	
		Fidonet 2:332/407	
	" "	+39 51 432904 ZyXEL 19.2k	
5		Fidonet 2:332/417	
NETHERLANDS:			
	YNY Koos z'n Doos	+31-3402-36647	
NEW ZAALAND:			
	YYY Advanced Systems	+64-9-379-3365 ZyXEL 16.8k	
5	Auckland Singet node number(s) 28:100/20 / 28:1000/201		
	INTLnet node number(s) 58:700/30 / 58:7200/30		
	Fidonet 3:772/360.10 (Just a pont sorry) 24Hrs Freqs etc.		
NORWAY:			
	Thunderball Cave	472567018	
?		RIME ->CAVE ?	
NETHERLANDS:			
	YNY DownTown BBS Lelystad	+31-3200-48852	14.4k
5	Lelystad	Fido 2:512/155, UUCP	
	YYY MUGNET Intl-Cistron BBS	+31-1720-42580	38.4k
4-5	Alphen a/d Rijn	UUCP	
SINGAPORE:			
	YYY The Controversy	(65)560-6040	14.4k V.32bis/HST
2-4		Fidonet 6:600/201	
SOUTH AFRICA:			
	NYN Pats System	+27-12-333-2049	14.4k v.32bis/HST
3	Pretoria	Fidonet 5:71-1/36	
SWEDEN			
	?Y? Gunship BBS	+46-31-693306	14.4k HST DS
?	Gothenburg		
SWITZERLAND:			
	NNY Atlantis	+41-1-492-8711	14.4k
4		ILINK	
	YYY Baboon BBS	+41-62-511726	19.2k
?		2:301/580 /581	
UNITED KINGDOM:			
	NYN The Purple Tentacle	+44-734-590990	HST/V32bis
4	Reading	Fidonet 2:252/305	

	A6 BBS	+44-582-460273	14.4k
?	Herts	Fidonet 2:440/111	
	YYY On The Beach	+44-273-600996	9600 HST
4	Brighton	Fidonet 2:441/122	

Commercial Service's Outside the US:

UNITED KINGDOM:

NNN Compulink Info eXchange 081-390-1255 v.32bis

5

4.2.11 What are File Requests?

FidoNet BBS's with the right type's of front-end mailer's can call other Fido BBS's and request their front-end mailer to send them files that they want. All this can be done automatically. File Requests (freqs) are basically the FidoNet equivalent to UUCP.

5 Installation and Common Problems

This FAQ section should be kept up-to-date, and should be the most *reliable* source for installation info. Please mail any corrections or changes to this section's coordinator, Matt Welsh (mdw@tc.cornell.edu). Last update April 1993.

5.1 Where to start

5.1.1 I want to install Linux on my machine. Where do I start?

The first thing you should do is read through the various introductory files, and ESPECIALLY the FAQ (this file). Especially this section :). A lot of effort has been done on these intro files, but note that some of them conflict with each other and cover older versions of Linux. When in doubt consult this file.

These files are all found on [tsx-11.mit.edu: /pub/linux/docs...](http://tsx-11.mit.edu/pub/linux/docs...)

FAQ The Linux Frequently Asked Questions list It's sitting in your hands now. This section is probably the best place to start to get the most up-to-date Linux installation information.

Info-Sheet Linux Information Sheet, by J. Winstead This is J. Winstead's general information sheet on Linux, and along with this FAQ is a good place to start for info.

README.kernel Kernel compilation README file, by L. Wirzenius This is the README notes for recompiling the Linux kernel from the sources. You don't need it unless you're planning to upgrade your kernel by compiling it yourself.

Others Every *release* of Linux (such as MCC-interim, SLS, HLU's disks, etc., see below) has its own up-to-date README files and docs that explain how to install that release. This FAQ section summarizes, but for more info on how to install Linux, read the READMEs and docs for the release that you choose.

Old docs There are a number of obsolete, old docs lying around. Most of these tell how to install Linux from the old boot/root disk combo. I DO NOT SUGGEST that you use these docs unless you know what you're doing — the best thing for beginners to do is read this FAQ and install the SLS release

(using the docs and READMEs for the SLS release). These old docs are things like `install.txt`, `guide.txt`, `RELNOTES`, `CHANGES`, and so on, and are all geared towards old versions of the boot/rootdisk. They are **not** relevant to current versions of Linux.

5.1.2 Is there some kind of limit on how large my Linux partitions and/or filesystems can be?

There's no limit on partition size (just the size of your drive), or filesystem size. The Minix filesystem type has a limit of 64 megs per filesystem. However, most folks use the Extended 2 filesystem or the Xia filesystem, which have limits of 4 terabytes per partition (virtually infinite).

See section 5.3 below on creating partitions and filesystems for more info.

5.1.3 How can I be sure I won't be writing over anything important when installing Linux? I have to use DOS on my machine, and I don't want to lose any files.

Back up everything. Just in case. As long as your DOS and Linux partitions don't overlap you should be okay.

5.2 Linux Packages

This section contains information about *some* of the current Linux packages available. The information related to the MCC package were written by A.V. Le Blanc.

5.2.1 Does there exist a way to get all (or nearly all) of the Linux stuff?

Yes. To install Linux, you're going to want to choose one of the "releases" of Linux, all of which have a different method of installation and set up. Each release also has its own README and installation docs, which you'll want to read first. But I'll summarize here. The major releases are:

The *Softlanding Linux System Release* Also known as the "SLS" release, consists of 24 disks for Linux and 10 for X11. The nice thing about this release is that you can pick and choose which disks and packages you want to install. The first disk must be "rawritten" (using `rawrite.exe`) on a floppy, and the rest of the images are put onto DOS format floppies. Contains all of the stuff in the MCC release plus a lot more (like mail/uucp), and includes X11 on the last 10 disks. Is also very easy to install for newcomers.

This is the release that I suggest everyone new to Linux should install. It's the most complete and up-to-date package. HLU's disks, below, are good for upgrading, and (unfortunately), the MCC-interim is quite dated at this point. If you install SLS you'll save yourself a lot of trouble.

The *TAMU (Texas A&M University) Linux Release* This release is supposedly like the SLS release, but has some different softs and a different installation procedure. From Dave Safford, "The installation procedure is the main difference from SLS. A single boot diskette is used, and it boots directly into an automated installation program." This installation program asks a few questions about the desired configuration, and sets up everything, including your filesystems, booting from the hard drive with LILO (see section 5.3.8 below), and a simplified X-Windows configuration.

This humble author has never installed the TAMU release but I've heard good things about it.

H.J. Lu's bootable rootdisk Release This is a release of the Linux kernel and basic binaries on a single floppy. It, along with HLU's *gccdisk*, *libdisk*, and so on, are good for upgrading or installing basic Linux system by hand. It's not recommended for newcomers, because there's no real install script, it's mostly meant as an upgrade of the basic system software. Beginners should install SLS or MCC-interim (see below) instead. The images and docs are found at `tsx-11.mit.edu: /pub/linux`.

The MCC Interim Release (Manchester Computing Centre) *Release* This is the fabled "MCC-Interim" Linux release, which was originally the de facto standard Linux distribution. After 0.97.2, it was dropped in favor of SLS, but as of 0.99.8 a new MCC-Interim has been released. Like SLS and TAMU, it has an easy-to-install, complete set of Linux software, including networking and X11.

Others but OUTDATED There are other releases and distributions of the Linux software, such as the "MJ" releas. There is also an older "boot/root" disk combo (0.98.1) which is like HLU's bootable rootdisk, above, but it's no longer supported (as far as I know).

The "MJ" release, according to Martin Junius, is no longer maintained. The last version was 0.97.1.

5.2.2 Where can I get these releases of Linux?

The "bootdisk/rootdisk" combo can be found at `tsx-11.mit.edu: /pub/linux/INSTALL`, as well as at many mirror sites (which is true for ALL of this software).

H.J. Lu's "bootable rootdisk" Release is found at `tsx-11.mit.edu: /pub/linux/packages/GCC`, as well as the mirror sites.

The MCC Interim Release is at `ftp.mcc.ac.uk: /pub/linux/mcc-interim`. It's also mirrored to `nic.funet.fi: /pub/OS/Linux/images/mcc-interim`, and `tsx-11.mit.edu: /pub/linux/mirrors/mcc-interim`.

The MJ release (discussed more below) can be found at `ftp.dfv.rwth-aachen.de: /pub/linux/mj-bin`.

The SLS Release is at `tsx-11.mit.edu: /pub/linux/packages/SLS`.

The sc.tamu versions can be found at `sc.tamu.edu` in `/pub/free_unix`.

5.2.3 What files do I need to get for the SLS release?

This is part of the FAQ written for by Peter MacDonald.

Basically all of the releases are alike. You need to get the DOS program `rawrite.exe` (or `rawrite2.exe`). This programs take a binary file and write it, block by block, to a blank formatted floppy. This is the way to take a Linux floppy image and put it onto a disk from DOS. `rawrite` can be found in `tsx-11.mit.edu: /pub/linux/dos_utils`.

For the SLS release, you need to get the files:

- `SLS.README` and `SLS.FAQ`. These files contain further information on the current release of SLS. You should read them in addition to this FAQ.

- a1.3 or a1.5, depending on what type of floppy you boot from on your system. a1.3 is for high-density 3.5" floppies, and a1.5 for high-density 5.25" floppies. This file is the image of the SLS "a1" disk. You must use **rawrite.exe** (above) to rawrite the file to a high-density **formatted** MS-DOS floppy. (Once the file is rawritten, the floppy will no longer be recognized by DOS).
- The files in the a2, a3, and a4 directories. These files must be copied to MS-DOS formatted floppies (using the MS-DOS **copy** command). For example, the files in the a2 directory are copied to a floppy labeled 'a2', and so on. These files make up the "a" disk set. They are required. The rest of the files below are optional.
- The files under the directories b1-bN, c1-cN, x1-xN, t1-tN, and so on. These files are copied to MS-DOS floppies just like a2-a4 above. (e.g. the files in the b2 directory are copied to a floppy labeled 'b2').

These floppies contain optional software for SLS:

a1-aN	The minimal base system
b1-bN	Base system extras, like man pages, emacs etc.
c1-cN	The compiler(s), gcc, g++, p2c, f2c etc.
i1-iN	Interviews (doc and idraw)
g1	german keyboard
s1-sN	sources for some of these utilities
t1-tN	T _E X(document processing)
x1-xN	The X-windows distribution

I suggest getting at least the a, b, and c disk series. If you want to run X-Windows, get the x series as well. When you're making these floppies, a2-a4 and so on must all be the same kind of floppy (i.e. either 3.5" or 5.25"). The a1 disk must be the floppy type that you boot from. The rest of the disks don't have to be the same type as the a1 disk.

5.2.4 Any other information about SLS?

This comes from the file **SLS.FAQ**, by Peter MacDonald. Please get the most recent version and read it: [sunsite.unc.edu: /pub/Linux/SLS/SLS.FAQ](http://sunsite.unc.edu/pub/Linux/SLS/SLS.FAQ).

SLS (Softlanding Linux System) Copywrite 1992, Softlanding Software.
which is NOT just an image dump of someones Unix system.

Also note that in the interest of preventing ftp storms, the version of SLS that appears on the internet, is not quite the same as the version distributed by Softlanding. Softlanding regularly gets a whole new version which has the changes integrated. But the updates to the Internet version are tailored to minimize the amount that has to be downloaded to become current. That is why bugs manage to creep in on me. I am not installing and testing the Internet version, although, functionally, it should be quite close to the Softlanding one.

So, why am I telling you this? After the next period of stability (few changes to SLS), I will be uploading the Softlanding version of SLS to tsx-11.mit.edu.

This distribution is freely available if you have internet access, or an obliging friend with access to it.

The purposes of the SLS are the following:

1. provide an initial installation program (for the queasy).
2. utilities compiled to use minimal disk space.

3. provide a reasonably complete/integrated UNIX system.
4. provide a means to install and uninstall packages.
5. permit partial installations for small disk configs.
6. add a menu driven, extensible system administration.
7. take the hassle out of collecting and setting up a system.
8. give non internet users access to Linux.
9. provide a distribution that can be easily updated.

SLS is a binary mostly distribution (except for the kernel), and is broken into multiple parts, or series, each of which is denoted by a letter followed by the disk number as follows:

a1-aN	The minimal base system
b1-bN	Base system extras, like man pages, emacs etc.
c1-cN	The compiler(s), gcc/g++/p2c/f2c
x1-xN	The X-windows distribution
i1-iN	Interviews (doc and idraw)
t1-tN	T _E X(document processing)
g1	german keyboard

This scheme allows new disks to be added to the distribution without changing the disk numbering. Also, the `sysinstall` program doesn't have to be changed when new disks are added as the last disk is marked by the presence of the file `install.end`. And when interviews is added, say as a new series `i`, it can be installed with:

```
sysinstall -series i
```

Highlights of the base are: `gcc/g++`, `emacs`, `kermit`, `elm/mail/uucp`, `gdb`, `sc` (spreadsheet), man pages, `groff`, `elvis`, `zip/zoo/lh` and `menu`. Highlights of X are: X, programmers libs, 75 dpi fonts, games (`spider`, `tetris`, `xvier`, `chess`, `othello`, `xeyes`, etc) and utilities like `xmag`, `xmenu`, `xcolormap` and `ghostscript`. Approximate usage is as follows:

Tiny base system:	9 MB	(Series 'a')
Main base system:	25 MB	(Series 'a', 'b' and 'c')
Main base system + X11:	45 MB	(Series 'a', 'b', 'c' and 'x')

Please read the file `COPYING` which outlines the GNU copying restrictions. The linux kernel is copywrite Linux B. Torvalds. Various other copywrites apply, but the upshot is that you may do whatever you like with SLS, except restrict others in any way from doing likewise, and you must leave all copywrites intact, and you can not misrepresent or take credit for others work.

Availability

SLS is available from the address:

Softlanding Software
 910 Lodge Ave.
 Victoria, B.C., Canada
 V8X-3A8
 (604) 360-0188

More details about SLS can be asked to pmacdona@sanjuan.uvic.ca

5.2.5 What should I do to install the SLS release?

Basically all of the releases are alike. You need to get the DOS program `rawrite.exe` (or `rawrite2.exe`). This program takes a binary file and writes it, block by block, to a blank formatted floppy. This is the way to take a Linux floppy image and put it onto a disk from DOS.

For the SLS release, you need to get the files in the `a1`, `a2`, `a3`, and `a4` directories at least. You also need the `README`'s there. Basically, you use `rawrite` to copy the `a1` and `a2` images onto floppies. Or, you can use `dd` on your UNIX system to transfer the files directly to floppies (assuming you have a floppy drive on your UNIX system) in place of `rawrite`.

The rest of the files all go onto DOS-formatted floppies, just using DOS copy. The SLS release is nice because it reads the DOS floppies during installation, so you don't have to `rawrite` all of those disks.

Once you have these 4 disks, you're ready to go. You can also get all of the files in the `b1-b7` directories (for extras, like man pages and emacs), `c1-c4` directories (for the GCC compiler and libraries), the `x1-x8` directories (for Xwindows), and so on, but they're all optional. I suggest at least getting the files in the `a`, `b`, and `c` directories. Remember that only the `a1` and `a2` images need to be rawritten, ALL of the other files just do onto the floppies in DOS format.

First you boot the `a1` disk. It will load up Linux, creating a ram-disk, and will ask you to put in the `a2` disk at some point. You'll be left with a prompt... from here you want to look at the SLS `README` file on-line (it tells you how). Then you'll run `fdisk` to create your Linux partition(s), and then reboot the `a1` disk again (putting in the `a2` disk, again, when asked). Then you'll run `mkfs` and `mkswap` to make your filesystems and swap space. Then you're ready to install the software - type `doinstall /dev/PART`

where `PART` is the main Linux partition you created with `fdisk`, above. From here on it's pretty automatic - you simply flip disks while it installs the software. First it will ask you how much software you plan to install - just choose the correct option depending on which sets of disks you got.

You will need to have a blank DOS-formatted floppy on hand. The installation procedure makes a Linux kernel boot disk out of it.

5.2.6 When installing SLS, the installation script creates a *boot floppy* for me to boot Linux with. How can I make these myself (or fix problems with them)?

The SLS install scripts just copy the kernel image to a blank floppy for you to boot Linux with. When SLS is installed, a copy of the kernel is stored in the file `/Image` on your hard drive: since you use the kernel on your boot floppy to boot, this file isn't used during the bootup process (unless you install LILO).

To make a boot floppy yourself, you must do two things:

1. Run `rdev` (a.k.a. `rootdev`) on the kernel in the file `/Image` to set the root partition it uses. For example, if your root partition is `/dev/hda3`, run the command

```
rdev /Image /dev/hda3
```

2. Copy the kernel to a new floppy. You may need to DOS format the floppy first (who knows why). Then run the command

```
cp /Image /dev/fd0
```

if the floppy is in `/dev/fd0` (the first floppy drive). You can use a command like

```
dd if=/Image of=/dev/fd0 bs=16k
```

as well; they should accomplish the same thing.

Many folks have problems with their SLS boot disks made in this manner because they forgot to run `rdev` on the kernel image before copying it to the floppy.

5.2.7 Where do I find the current releases of programs?

New releases and new programs are usually announced in `comp.os.linux` and/or the mailing lists. Other than that, you more or less have to check the FTP sites.

Linus' `.plan` contains some information about the current kernel. Finger Linus Torvalds `finger torvalds@klaava.helsinki.fi`. (It is often quite long!)

Please try to use the current releases, if possible, at least of the kernel and compiler. If you have a problem, please first make sure you have a current release, it often helps. Don't use alpha test releases unless you *like* problems.

5.2.8 How do I get and install H.J. Lu's "bootable rootdisk" release?

Basically follow the directions above for the 2-disk "combo" release (i.e. get the images, rawrite them to high-density DOS formatted floppies, boot off the boot disk). This release doesn't contain user-friendly installation scripts; it's meant for those who can install the images from the disks by hand (the `READMEs` contain instructions on doing this).

5.3 Some common Problems

5.3.1 What filetype is the extension `.z`? What about `.taz`, `.tpz`, and `.tgz`? I see these files on the archives but I don't know how to unpack them.

Here's a list of common filename extensions on the archives.

`.tar` Tape ARchive, created with `tar`. To unpack type

```
tar xvf Archive.tar
```

If you can't access the files, because you are not the owner, try

```
tar xvof Archive.tar
```

Or, you can do

```
tar tvf Archive.tar
```

to get an index listing of the tarfile.

`.Z` This file was compressed with `compress`. You can uncompress it with

```
uncompress Archive.Z
```

or

```
zcat Archive > file
```

.z **gzip** is now used by many archive sites instead of **compress**; if you don't have **gzip** on your system, get it! To uncompress one of these files use

```
gzip -d foo.z
```

.tar.Z This archive is compressed, so you have to uncompress it before unpacking. It's faster and uses less disk space if you use

```
zcat Archive.tar | tar -xvf -
```

.TZ, .taz These filenames are made for filesystems that don't allow filenames longer than 14 characters. They are renamed **.tar.Z** archives. To unpack them you can't use **zcat** or **uncompress** the normally way because they will add a **.Z** to the filename. Use

```
zcat < Archive.taz | tar -xvf -
```

instead or

```
tar xvfz foo.taz
```

to unpack it (some versions of **tar** don't have the **z** option).

.shar This file is a shell-archive. The archive is a shell-script which will unpack the files. To extract read the header first. Often you have to delete parts of the header before executing with

```
sh Archive.shar
```

.uue This archive is converted into 7-bit-code with **uuencode**. The files are bigger than before, but you can transfer them via mail. To restore the files, use

```
uudecode Archive.uue
```

.diff These files are no real archives. They include information about differences to an old version of a file. The program **patch** uses this information.

.F This archive is compressed by a tool from the BSD-world called **freeze**.

.zip An archive format imported from the DOS-world. Nowadays it's interesting for UNIX users, too, because the Free Software Foundation had made a decision to pack their files with this format. That's why there is a GNU-Zip program.

.tpz, .tgz **gzipped tar** file. If you have **gzip**, **zcat** is linked to it, so you can do

```
zcat foo.tpz | tar xvf -
```

to unpack it. `.tpz` is the old extension; all gzipped tar files should now end in `.tgz` instead.

The SLS distribution uses gzipped tar files (`.tgz`).

5.3.2 How do I make partitions and/or filesystems for Linux?

For most Linux installations (such as SLS) you'll need at least two partitions: one for swap space (used as virtual memory) and another for your *root filesystem* (that is, the actual Linux software itself). You can also make separate partitions for your `/usr` filesystem, etc. (however, the SLS `doinstall` program, at this point, mounts your root filesystem for you before installing, thus, you can't have a separate `/usr` filesystem to start out with), if you have UNIX experience, you'll be able to figure out how to divide software between multiple partitions.

For most installations simply having one root filesystem and one swap partition is the best way to go.

But that's beside the point. First thing you need to do is resize the existing partitions on your drive (if any) to make space for Linux; for example, if you have a DOS partition taking up all of your drive, you need to use `FDISK` under MS-DOS to delete it and recreate it with a smaller size. Of course, in so doing you'll lose everything on that DOS partition — just back it up first and reinstall after you've recreated and reformatted the partition. That's life! :) (no longer as FIPS, a non-destructive disk partitioner, is ALPHA version [Joey])

Now you can boot up Linux (say, from SLS, or from the boot/root disks). From there you use the command `fdisk` to create partitions for Linux:

```
fdisk <drive>
```

where `<drive>` is the name of the drive which you're installing Linux on. Valid drive names are:

```
/dev/hda : first IDE drive
```

```
/dev/hdb : second IDE drive
```

```
/dev/sda : first SCSI drive
```

```
/dev/sdb : second SCSI drive
```

For example, to run `fdisk` on the first SCSI drive in your system, use the command `fdisk /dev/sda`. If you don't specify a drive name for `fdisk`, it will assume `/dev/hda`.

The `fdisk` command `m` prints a menu. The `n` command is used to create a new partition. `w` writes the partition table and exits `fdisk`; `q` quits without saving changes. The `p` command displays the current partition table.

When creating partitions with `fdisk`, you may get an error message to the effect of **Warning: Linux cannot currently use xxxxx blocks of this partition**. This simply means that the Minix filesystem type (which is the default) is limited to filesystems which are 64 megs in size. You can ignore this warning, because when creating your filesystems, you can use one of the other filesystem types (`ext`, `ext2`, or `xia`) which have a size limit of 4 terabytes.

Just keep in mind that if you create a partition larger than 64 megabytes then you cannot use the Minix filesystem type ('mkfs') when creating the filesystem (see below). You must use the `ext`, `ext2`, or `xia` filesystem type instead.

Extended filesystem does NOT equal *extended partition*. An extended partition acts as a *container* for logical partitions, which is nice because you can have many logical partitions inside the extended partition, and only use up one of your four available primary partitions on the drive. (You can have only one extended partition per drive). Extended partition can't hold data on their own: you need to create logical partitions on top of it first. They are numbered `/dev/hda5`, `/dev/hda6`, and so on.

The extended filesystem on the other hand has nothing to do with extended partitions. It's just the name of another filesystem used by Linux which lets you have larger filesystem sizes and other features.

The TYPE of the partition (in the rightmost field when you use the `fdisk "p"` command) does not matter to Linux. However, it does matter to other systems such as OS/2. Just make sure that your Linux partitions have a type such as "Linux/MINIX" to distinguish it from other systems. (i.e. It doesn't matter if you set the type of the swap partition to "Linux swap". All of your Linux partitions can use the default type "Linux/MINIX" which is fine).

WARNING: DO NOT use Linux `fdisk` to manipulate partitions for operating systems other than Linux. This means don't use Linux `fdisk` to delete or create MS-DOS partitions, for example. Use MS-DOS `fdisk` for this instead. Linux `fdisk` doesn't properly know how to create partitions for MS-DOS and other operating systems.

When creating partitions, write down the names and sizes of the partitions you create. You need this information later on. Use the `fdisk p` command to display this information.

Once you've got your partitions created, you need to reboot the system so that changes to your partition table are picked up. Then you need to **format** the partitions (i.e. create filesystems on them). For your swap partition, use the command `mkswap <partition> <size>`. For example, if you have a swap partition on `/dev/hda2` with a size of 8208 blocks (about 8 megs), use the command

```
mkswap /dev/hda2 8208"
swapon /dev/hda2
```

This will make the swap space and enable it. Swap space is used as virtual memory... if you have 4 megs of RAM in your machine, and a 16 meg swap partition, you have a total of $4 + 16 = 20$ megs of RAM available for Linux.

To make your filesystems, use the correct version of the `mkfs` program, depending on the type of filesystem you wish to use:

Filesystem type	Command to create filesystem
Minix	<code>mkfs -c <partition> <size></code>
Extended (old)	<code>mkefs -c <partition> <size></code>
Second Extended	<code>mke2fs -c <partition> <size></code>
Xia filesystem	<code>mkxfs -c <partition> <size></code>

The Minix and Extended filesystem types are older, and while more robust, are slower and may eventually be phased out. The most popular filesystem type used right now is the Second Extended fs type ('mke2fs'). The Minix filesystem type is limited to 64 megs (and 14-character-long filenames). It's suggested that you use either `mke2fs` or `mkxfs`.

For example, to create a 61000-block Second Extended filesystem on `/dev/hda3`, use the command

```
mke2fs -c /dev/hda3 61000
```

The `-c` option checks for bad blocks on the filesystem. Apparently `mkfs` and `mkefs` do not actually use the bad block information correctly (see Q/A below).

Now you're ready to install the software. For SLS this means just using the `doinstall` command (see section 5.2 above). For other installations you may need to *mount* the filesystems first. Just see the relevant **READMEs** for information.

5.3.3 Why does `fdisk` say Linux cannot currently use XXXX sectors of this partition?

See previous question. `fdisk` is an older program which expected all filesystems to be Minix fs, which limited filesystems to 64 megs. You can ignore this warning because the `extfs`, `xiafs`, and `ext2fs` don't have this limit.

5.3.4 What does the message `MINIX-fs: Magic match failed on bootup mean?`

Basically, this means that you're trying to mount a non-Minix filesystem as a Minix filesystem, and `mount` is croaking on it because the type is wrong. If you use a root filesystem type other than Minix fs (i.e. if you use the `extfs`, `ext2fs`, or `xiafs`) you'll probably see this message: the kernel tries to mount root as Minix, then `extfs`, then `ext2fs`, and so on... every time it fails for one type it tries the next type.

In most situations, this can be ignored, if the kernel is able to mount your root partition as any one of its known types. However, if you have the root device set to the wrong partition, for example, then the kernel should hang at this point and you'll have to use `rdev` to set it correctly.

5.3.5 Linux `mkfs` doesn't accept the size I give the device, although I double-checked with `fdisk`, and it's correct.

Be sure you give the size in Linux BLOCKS (1024 bytes), not sectors. Also make sure that you have the right partition: partitions are numbered `/dev/hda1`, `/dev/hda2`, and so on (and `/dev/hdb1`, `/dev/hdb2` for the second hard drive)... DON'T use `/dev/hda` or `/dev/hdb` as they correspond to the entire disk, not just single partitions.

Also remember that SCSI drives use `/dev/sda1`, `/dev/sda2`, ... and `/dev/sdb1`, `/dev/sdb2`, and so on for their partitions.

5.3.6 How can I get `mkfs/mkefs` to check for bad blocks?

Unfortunately the `-c` option on `mk(e)fs` does not work; it cannot detect bad blocks on the drive, and thus if you create a filesystem over a part of your hard drive with bad blocks, things will eventually go wrong. So you need to generate a bad block list (in a file) and use the `-l` option on `mk(e)fs` so it will flag those blocks when making a filesystem.

`mke2fs` and `mkxfs` correctly flag bad blocks, so this only applies to `mkfs` and `mkefs`. This is also only needed for older RLL and MFM drives. SCSI and IDE drives have bad block logic on-board.

How to generate a bad block list? Since this is only needed for older drives and older filesystem types (i.e. Minix and `extfs`), just mail `mdw@tc.cornell.edu` or `kircher@neuro.tuwien.ac.at` if you need instructions on this procedure. Your best bet is to use `mke2fs` or `mkxfs` anyway.

5.3.7 Generating a bad block list for `mk(e)fs`

This item was written by Gerhard Kirchner (kirchner@neuro.tuwien.ac.at) on Nov 11, 1992

0. Introduction

First of all: the `-c` option (check bad blocks) of `mk(e)fs` does not work (the code used cannot detect any bad blocks). Fortunately there is another option `-l <file>` to tell `mk(e)fs` where the bad blocks are. SCSI and IDE drive users do not need any bad block management — the drive logic does it for them. However, users of old MFM or RLL drives do need it.

1. Where are my bad blocks?

Every MFM or RLL drive is tested by the manufacturer and the locations of bad blocks are usually printed directly on the drive case. If you cannot find any bad block information on your drive, you can use the common disk test software (Norton,...) to check your drive and obtain a defect list.

Locations of bad blocks are given in terms of cylinder and head coordinates. Cylinder and head together specify a certain track. Some manufacturers are more specific about the location of the bad spot within the track but if you know how to use this information you probably do not need to read this guide. So if we know the track where the bad spot sits, we declare the entire track as bad (even experts do this).

What we have now is a table of the form

CYLINDER	HEAD
...	...
...	...
...	...

2. Calculating all bad sectors

Each sector on the drive can be addressed by three coordinates:

Cylinder	=	[0..CYLINDERS-1]
Head	=	[0..HEADS-1]
Sector	=	[0..SECTORS-1]

where

CYLINDERS		total number of cylinders
HEADS		total number of heads
SECTORS		number of sectors per track

The absolute address of a sector on the disk is then calculated according to the formula

$$Abssector = HEADS * SECTORS * Cylinder + SECTORS * Head + Sector$$

We can now calculate the absolute addresses of all sectors of every bad track on our disk according to

```
for each bad track given by Cylinder and Head do
begin
  for Sector=0 to SECTORS-1 do
```

```

begin
  Abssector = HEADS*SECTORS*Cylinder + SECTORS*Head + Sector
end
end

```

obtaining a list of all bad sectors.

3. Converting absolute sectors to partition relative sectors

Partitions are like separate disks, that means that sector counting starts with 0 for each partition. So what we have to do next is to generate a bad sector list for each partition we want to use for linux. To do that, we must know where each partition starts and ends. We can get this information from `fdisk`. When we start `fdisk` and type `p` to view the partition table, the begin, start and end information will be given in terms of cylinders. When we type `u` to toggle the units, and then again `p` we get what we need: Start and End of each partition in terms of absolute sectors. Here is what I see on my system (I use my second disk for linux):

```

# fdisk /dev/hdb

Command (m for help): p

Disk /dev/hdb: 8 heads, 17 sectors, 1024 cylinders
Units = cylinders of 136 * 512 bytes

   Device Boot      Begin   Start    End  Blocks  Id System
/dev/hdb1          1         1    963   65483+  81 Linux/MINIX
/dev/hdb2          964        964   1024    4148   82 Linux swap

Command (m for help): u
Changing display/entry units to sectors

Command (m for help): p

Disk /dev/hdb: 8 heads, 17 sectors, 1024 cylinders
Units = sectors of 1 * 512 bytes

   Device Boot  Begin   Start    End    Blocks  Id System
/dev/hdb1          1         1  130967   65483+  81 Linux/MINIX
/dev/hdb2       130968  130968  139263    4148   82 Linux swap

```

Sector counting (as counting always should do :-)) starts with 0. The first partition begins with sector one, as sector 0 is always the boot sector.

This is how we calculate the addresses for one partition:

- (a) From the list obtained in section 2, cancel all addresses that are not in the range [Start..End] (including limits) shown by `fdisk`
- (b) Subtract Start from each remaining entry

4. Converting to blocks

A disk sector has a length of 512 bytes (this is the usual size BIOS/DOS can handle). Linux groups two sectors to an entity called block. A block therefore consists of two physical sectors and has a size of 1024 bytes (I read

somewhere that larger blocks will be supported in the future). The `mk(e)fs` program wants to know the addresses of bad blocks, not sectors. So we convert our partition relative sector addresses to partition relative block addresses by dividing by two and taking the integer part. Doing that we certainly get a lot of duplicate addresses which we have to get rid of, so we simply delete redundant entries.

What we have now is a list of bad blocks for each partition. `mk(e)fs` likes to get this information from a file, one address per line, one file for each partition.

5. Automating the process

Doing all the calculations by hand is tedious and error prone. So we may decide to automate the whole thing. We can do everything in dos when we use the proper utilities (see later)

- The Input File

We need a single file containing the coordinates of all bad tracks of one entire disk. Each line of the file consists of two fields, the first of which is the cylinder number and the second is the head number. These are the first few lines of the file for my second disk:

```
48 0
105 4
150 2
224 1
380 2
427 6
435 1
```

- An AWK script for doing all the calculations

In the following we present an awk script that does all the work for us. We just have to plug in the correct values of the partition and drive parameters and off we go. I use `gawk211.zip` and `sort03.arc` which I downloaded from a simtel20 mirror. Sort is used to remove the duplicates.

```
#---the awk script starts here
# generates linux bad blocks file (starting count with 0)
# for partition ONE
# for micropolis drive 8 heads, 17 sectors
# input file: cyl[0.. ] head[0.. ]\n
# stdout: abs blocks of corresponding partition\n
BEGIN { start = 1;
        end = 130967;
        sectors = 17;
        heads = 8;
        }
{ sec=$1*sectors*heads+$2*sectors;
  if (sec>=start && sec<=end)
    for (i=0; i<sectors; i++) print int((sec-start+i)/2) | "sort -mu"
}
#---the awk script ends here
```

Suppose that our input file is named `badtrack.lst` and that the awk file is `part1.awk`, then we can generate a bad block list for `mk(e)fs` by typing the command

```
gawk -f part1.awk badtrack.lst > badblk1.lst
```

This is the file we need.

6. Making the file system

What we must do now is to have this file accessible for linux during installation. There are several ways to do this. You can either mount your DOS partition and copy the file over to Linux, or use `mttools` (as in `mcopy c:badblk1.lst /user/badblocks`).

Now we can make the file system. In this case the bad blocks file is in `/user/badblocks`, and the partition we're making a filesystem on is `/dev/hdb1`. We're making an extended filesystem with a size of 65483 blocks:

```
mkefs -l /user/badblocks /dev/hdb1 65483
```

Thats it!

7. Bugs

I only tried `mkfs` but I'm rather convinced that it works with `mkefs` as well.

5.3.8 How can I boot Linux off of my hard drive?

You need to install the LILO program which changes the boot sector of your hard drive to allow you to choose between a DOS or a Linux partition to boot from. These programs are provided with most major releases, or you can get them seperately from one of the FTP sites.

As of LILO version 8 there is a *quick install* script available which should make LILO installation quick and easy. See Section VII for LILO information. You can also see the *LILO QUICKSTART GUIDE* by Matt Welsh which should be available on `sunsite.unc.edu: /pub/Linux/docs`.

5.3.9 Cripes! I tried to install LILO, but screwed up somehow, and now can't boot anything from the hard drive. How can I fix this?

Simple. First you need an MS-DOS 5.0 bootable floppy with `FDISK.EXE` on it. Boot it and run the command

```
FDISK /MBR
```

which will (hopefully) restore your hard drive's master boot record to a standard MS-DOS boot record. Now you can go back and reinstall LILO. :)

5.3.10 When installing SLS, I get the error You may have inserted the wrong disk when putting in the next disk in a series. What's going on?

Each disk has a small file on it which contains the name of the disk. For example, the SLS a3 disk has a file on it called `diska3`. If it doesn't exist, or is named something else (like `diska3.z`), then just create it/rename it.

Also, the last disk in a series (i.e. the a4 disk, b5 disk, etc.) has a file on it called `install.end`. You need this file as well. These files are used by the SLS `doinstall` program to keep track of where it is.

5.3.11 When installing SLS, the installation script creates a *boot floppy* for me to boot Linux with. How can I make these myself (or fix problems with them?)

The SLS install scripts just copy the kernel image to a blank floppy for you to boot Linux with. When SLS is installed, a copy of the kernel is stored in the file `/Image` on your hard drive: since you use the kernel on your boot floppy to boot, this file isn't used during the bootup process (unless you install LILO).

To make a boot floppy yourself, you must do two things:

1. Run `rdev` (a.k.a. `rootdev`) on the kernel in the file `/Image` to set the root partition it uses. For example, if your root partition is `/dev/hda3`, run the command

```
rdev /Image /dev/hda3
```

2. Copy the kernel to a new floppy. You may need to DOS format the floppy first (who knows why). Then run the command

```
cp /Image /dev/fd0
```

if the floppy is in `/dev/fd0` (the first floppy drive). You can use a command like

```
dd if=/Image of=/dev/fd0 bs=16k
```

as well; they should accomplish the same thing.

Michael L. Kaufman (kaufman@delta.eecs.nwu.edu) says: Formatting the floppy lays down the track/sector information. If you have a completely unformatted floppy, `dd` can fail in interesting ways.

Many folks have problems with their SLS boot disks made in this manner because they forgot to run `rdev` on the kernel image before copying it to the floppy.

5.3.12 How can I set the default video mode used by Linux? Do I have to recompile the kernel to do this?

No, you don't have to recompile the kernel. Just use `rdev` with the `-v` switch to set the video mode in the kernel (either on your hard drive (if booting from LILO) or on your boot floppy). For example use if booting with a floppy disk.

```
rdev /dev/floppy
```

This tells you the default root device, such as

```
Root device /dev/hda1
```

If booting off the hard drive, then it must be done on the boot device whatever that may be (`/dev/hdax`, `/dev/hdbx`, `/dev/sdax`).

Note that it is not the `Image` file which needs to be changed, but the `BOOT` device. (I found out after much hassle and frustration)

For swap device use `-s`, for video mode `-v`, for ramdisk size `-r`.

```
rdev -v /dev/floppy
```

shows the current video mode. The choices are (can be found out from `rdev -h`):

```
-1 Normal VGA
-2 Extended
-3 Ask (on bootup)
 1 mode 1 (as selected by keypress on bootup)
 2 mode 2
...
 6 mode 6
```

To set bootup from hard disk (`/Image`) with video mode 4 do

```
rdev -v /Image 4
```

To reset to asking on boot, do

```
rdev -v /Image -3
```

Therefore, it is changeable/reversible, and does not necessarily require recompiling the kernel.

5.3.13 How else can I use rdev?

`rdev` is very handy and also is used to set the root and swap partitions, ramdisk size, and more, in a compiled kernel. It means you don't have to recompile the kernel to make these changes. Use `rdev -?` for a list of options.

5.3.14 When I login as non-root, I get tons of errors about shell-init: permission denied. Also, some things work as root but not as a normal user. What's the deal?

This is a really common problem which comes from not having permissions set right and a misunderstanding of some UNIX terms. Some installations won't have the file permissions set correctly on the various directories that normal users (i.e. non-root) will use. For example, if your user directories are in `/home`, then `/home` must be of mode `rwrx-xr-x`, or `chmod 755 /home`. Also, a home directory must be owned by the user who it belongs to (i.e. `/home/mdw` must be owned by `mdw`... just `chown mdw /home/mdw`, for example). Also, the permissions must be set correctly for `/` (the root directory). Here's a list of permissions that should work (although you can use other permissions, these are just suggestions that shouldn't cause trouble):

permissions	(chmod)	owner	file
<code>rwrx-xr-x</code>	(755)	root	<code>/</code>
<code>rwrx-xr-x</code>	(755)	root	<code>/home</code>
<code>rwrx-xr-x</code>	(755)	mdw	<code>/home/mdw</code>
<code>rwrx--r--</code>	(744)	mdw	<code>/home/mdw/.profile</code> (or other startup files)
<code>rwrx--r--</code>	(744)	mdw	<code>/home/mdw/foo</code> (normal files)

In most cases the group of the file doesn't matter, but in general most files are set to group `root` (except for binaries which are group `bin`, and so on) and user files are set to group `user` (or whatever group users are in).

In general you want directories that everyone can access to be `rwrx-xr-x` (chmod 755). Files that everyone can read are `rwrx--r--` (chmod 744), and programs that everyone can run are `rwrx-xr-x` (chmod 755).

Thus `/bin`, `/usr`, `/usr/bin`, `/etc`, and so on, should all be `rwrx-xr-x`. All binaries should be `rwrx-xr-x` (unless, of course, they're setuid programs). For UNIX newbies, setuid programs run under the user id of the owner, thus programs owned by root which have a permission of `rwsr-xr-x` (note the `s`) run as root, with

root's privileges. So before changing permissions on a program check to see if it's setuid first. To make a program setuid prepend a **4** to the permissions you give to **chmod**, i.e. **rwsr-xr-x** is **chmod 4755**.

Here's a list of common uses for permissions:

permission	effect on files	effect on directories
read	lets you look at a file, lets you run a shell script	lets you see what's in the directory with ls
write	lets you edit a file or copy over it	lets you delete a file from the directory ²
execute	lets you run a binary or shell script	lets you cd into the directory

5.3.15 I have the previous Linux version, how can I upgrade it?

If you've never done this before, get the kernel sources from your nearest FTP site (in a file named something like **linux-*.tar.Z**) and unpack them into **/usr/src/linux**. As of Linux-0.99 there is a script to autoconfigure your kernel, and you need at least GCC version 2.3.3. You unpack these sources in **/usr/src/linux**. Make sure you have the logical links for **/usr/include/linux** and **/usr/include/asm** as described below.

First run **make config** and answer the various questions. Then edit **/usr/src/linux/Makefile** to set the root partition, keyboard, etc. Finally, do a **make dep** (to set dependencies: VERY important!) and finally **make**. Assuming you have GCC installed correctly, the kernel should compile and you'll be left with a new **Image** which is your new kernel: if you boot from harddrive, copy the Image to wherever you told LILO to look for it, or if you boot from floppy dd the Image to a new floppy. Make sure that you run "rdev" on the Image to make it look for the correct partition for your root filesystem (if you specified this correctly in the Makefile you don't need to do this).

Make sure you read **/usr/src/linux/README**, which explains in detail what to do when recompiling the kernel.

If you **HAVE** done this before, you can just apply the source patches to your old sources and then recompile (i.e. you don't have to get the entire kernel sources all over again). Use the **patch** program to do this. Before you recompile the kernel do a 'make dep' to set dependencies and a **make clean**.

5.3.16 Where is /usr/include/linux and /usr/include/asm, or, Why won't the kernel compile correctly?

The files **/usr/include/linux** and **/usr/include/asm** are symbolic links to **/usr/src/linux/include/linux** and **/usr/src/linux/include/asm**, respectively. In other words, all of the Linux include files are actually under **/usr/src/linux/include**, but to access them you need two symbolic links in **/usr/include**:

```
/usr/include/linux -> /usr/src/linux/include/linux
/usr/include/asm -> /usr/src/linux/include/asm
```

To make these, run the commands

```
ln -sf /usr/src/linux/include/linux /usr/include/linux
ln -sf /usr/src/linux/include/asm /usr/include/asm
```

NOT the other way around. :) If you don't have these links then many compilations will surely fail.

5.3.17 How can I upgrade to the newest version of GCC?

See section 12 on GCC for more information, but essentially all you need to do is FTP to `sunsite.unc.edu` (or one of the other Linux FTP sites) and look in `/pub/Linux/GCC`. You'll see a number of `.tar.Z` files there (often abbreviated `.TZ` or `tpz`). The names change from time to time: you need to get the compiler (often in a file such as `gcc233.TZ`), header files, and library files. At this point, the compiler lives in `/usr/bin`, the compiler's setup files are under `/usr/lib/gcc-lib/i386-linux`, and the libraries (shared, jump table, and static) are all in `/usr/lib`. It's all very straightforward once you actually unpack the tar files and everything falls into place.

5.3.18 What's the deal with these things called *jump tables*?

There's more information in sections 12, 8 and 9, but to be brief: There are three kinds of libraries for Linux. As of `gcc v2.3.3` they all live in `/usr/lib`. The files are...

```
/usr/lib/*.a static (non-shared) libraries (use gcc -static ...)
/usr/lib/*.ca classic shared lib stubs (use gcc -nojump ...)
/usr/lib/*.sa jump table shared lib stubs (use gcc -jump ...)
```

There is an older form of the shared libraries, classic shared libraries, which is no longer widely used. Nonetheless, you may run across it from time to time.

When you compile a program, depending on the options you give `gcc` (`-jump` is the default) it will link it against one set of these libraries. The static libs contain all of the code and thus make your executable very big; no shared code is used. The classic shared libs are actually `stubs` which reference the shared code in

```
/lib/libc.so.VERSION (where VERSION is a number like '4.2').
/lib/libc.so.VERSION is a file which contains the actual code of the library, which
is accessed at runtime by your executable. The jump-table library stubs are also
shared libs, but they are built in such a way that you can upgrade the /lib/libc.so.
VERSION file without having to recompile the programs that use it.
```

So when we say *uses Jump Tables version 4.2* we mean it uses the actual library itself, `/lib/libc.so.4.2`. To compile programs that use jump tables version 4.2 you need the right version of `/usr/lib/*.sa` installed, but you don't need them to RUN programs that use jump tables.

However programs that are compiled to use, for example, jump table version 4.2 (in the file `/lib/libc.so.4.2`) won't work if you only have `libc.so.4.1` installed. They're only backwards-compatible. If you get errors about can't find `/lib/libc.so.4.1` then the executable you're using is looking at runtime for a jump table version that you don't have. Basically you're safe if you have the most recent version of the

`/lib/libc.so.VERSION` file installed (which is found on `sunsite.unc.edu: /pub-verb"/Linux/GCC`" and comes with the GCC stuff).

You should always have a symbolic link with the major version number of the library in `/lib`, because that's what's actually read. For instance, if you have `/lib/libc.so.4.2` installed, then you need the symbolic link

```
/lib/libc.so.4 -> /lib/libc.so.4.2
```

Make this with the command

```
ln -s /lib/libc.so.4.2 /lib/libc.so.4
```


5.3.19 How to upgrade jump tables? (Without hanging my system)

Be careful! If you upgrade your `/lib/libc.so.VERSION` file and either remove the symlink or delete the old `libc.so.VERSION` file before the new one is in place (and the link points to it), then more than likely all of your binaries like `cp`, `mv`, `ls` and `ln` use the library, so they'll all stop working once you kill the link or the library file. So to upgrade the file, copy the new version to `/lib` and switch the link in one step with a command such as

```
ln -sf /lib/libc.so.NEW_VERSION /lib/libc.so.4
```

where `NEW_VERSION` is the new version of the library you're installing. This will switch the link in one step and everything should work. You just can't copy over the old file because the old version is in use by `cp`, `mv`, and so on.

5.3.20 How can I be sure I won't be writing over anything important? I have to use DOS on my machine, and I don't want to lose any files.

Back up everything. Just in case. Then, write some easily recognizable pattern to the partition you have reserved for Linux, using some DOS tool. You can then use `cat /dev/hdX` under Linux to examine which of the partitions you used.

5.3.21 I just rebooted my machine, and now Linux dies with a panic: trying to free unused inode. What's going on?

You probably forgot to `sync` before rebooting, which stores on the disk physically the contents of the kernel buffers. You can either run `fsck` on the partition to TRY to correct the problem (it might fail), or `re-mkfs` and `re-install` the software on that partition.

The best way to shutdown your system is the `shutdown` command. To shutdown and reboot the system, use

```
shutdown -r now
```

or substitute a number of minutes in place of `now`. Leave off the `-r` switch if you just want to shutdown and not reboot.

You can now use the `reboot` command, which syncs all disks, shuts down the system and reboots it.

5.3.22 Can I use both OS/2 and Linux on my machine?

Yes! See the following two Q/A's about getting your OS/2 Boot Manager to work. But, be warned: **IF YOU USE OS/2, DO NOT USE LINUX'S fdisk TO CREATE LINUX PARTITIONS!!** The problem is with a bug/feature in OS/2's `fdisk` that tries to correct *errors* in partitions that it doesn't like... Linux partitions included. The solution: make your Linux partitions with OS/2's `fdisk`, then use Linux's `fdisk` to change the partition ID's to the right values (this is self-explanatory with Linux's `fdisk`).

If you made your Linux partitions with Linux's `fdisk`, and OS/2 sees them, it will think they have errors and end up trashing them.

5.3.23 I use OS/2's Boot Manager on my hard drive. How can I get it to recognize Linux?

To do this, install LILO on your Linux root partition, NOT on your hard drive's master boot record. The lilo command for this would be (if `/dev/hda3` is your Linux root partition, and your Linux kernel is in `/vmlinuz`):

```
/etc/lilo/lilo -c -b /dev/hda3 -v -v /vmlinuz
```

Then use OS/2's `fdisk` to add it to the Boot Manager.

Thanks to Thomas Brodt (brodttts@minnie.informatik.uni-stuttgart.de).

5.3.24 When I run Linux's fdisk it says OPUS for OS/2's Boot Manager partition. Is this right? What's OPUS?

It's correct. OPUS is BBS software that used partition type 0x0A long before OS/2.

5.3.25 Why doesn't the termcap entry for VT100 match the entry for a VT100 in a sun termcap?

I would like a complete termcap, but if there isn't one, what differences do I need to know about with regards to Linux if I have to write one myself. For example, I tried using the `h19` entry from the sun termcap and it didn't work.

That's because our vt100 emulation isn't completely done. It varies in small ways from an actual vt100. You would have to ask someone who worked on it what the actual differences are.

5.3.26 There is no include/types.h.

That's because it's called `sys/types.h`.

5.3.27 What are the differences between the different system inits? And will they work with the getty's and shadow password packages?

Init just does whatever it's told to do in `/etc/inittab`. So `getty's` or any `getty` for that matter will work with any init if `inittab` is set up correctly. As far as shadow passwords go, that's login's responsibility to get correct, so it doesn't matter what init you use.

5.3.28 I tried following the documentation for TCP/IP, but I don't know much more than the basics for networking. Currently I just want to telnet to myself. When I try, I get a general protection, then a seg fault, followed by: named[34]: No root nameservers for class 1.

Is this stuff part of the `ka9q` package or do I need to get that also?

Try reading the Net FAQ. It was updated recently, and should help a lot. I currently can Telnet to my Linux box from school through the term program, and had no problems setting up the networking stuff.

6 Some Classical Problems

6.1 Kernel

6.1.1 While running `du` I get Kernel panic: `free_inode: bit` already cleared. Also, `du` produces a `ENOENT` error for all the files, in certain of my directories. What's going on?

These are both consistent with a bad file-system. That's relatively easy to produce by not syncing before rebooting, as Linux usually has 1.5MB of buffer space held in memory (unless you have ≤ 4 MB RAM, in which case the buffers are only about 0.5MB). Also Linux doesn't do anything special about the bit-map blocks, and as they are used often, those are the thing most likely to be in memory. If you reboot, and they haven't been written to disk ...

Just do an `fsck` on the device, the `-a` flag might repair it otherwise, the only thing to do is to reinstall the filesystem from the Images.

A sync is done only every 30 seconds normally (standard UNIX practice), so do one by hand (some people think you should do 3 syncs after each other, but that's superstition), or by logging out from the startup-shell, which automatically syncs the system. Unmounting a filesystem also syncs it (but of course you can never unmount root).

Another (sad) possibility is that you have bad blocks on your disk. Not very probable, as they would have to be in the inode-tables, just a couple of blocks in size. Again there aren't programs available to read a disk for bad sectors and put them in some kind of *bad-sector-file*. On IDE drives this is no problem (bad sectors are automatically mapped away).

6.2 Drives

6.2.1 How can I partition my hard-drive to use Linux?

See section 5 of the FAQ on installation.

6.2.2 I heard something about repartition a hard disk without deleting everything on it, any clue?

It's not a program but a partition procedure which requires

1. a partitionning program
2. a sector editor

The procedure itself can be found (at least) in `digest#132 Vol2`.

6.2.3 What must I do to `mkfs` a floppy?

Blocks are of size 1K so 1.44 floppy is 1440 blocks. The floppy has to be formatted before this will work (e.g., `fdformat` can do this from within Linux).

6.2.4 I have some trouble with `tar`/`untar`.

The `tar` provided on .96 and latter is Pax (don't know for .98) which do not accept the `z` flag. You can download the GNU tar at `tsx-11` in `/pub/linux/binaries/usr.bin`.

6.2.5 I can do this as root but not as non-root, is it a bug?

Except for an early make utility, the problem is caused by an incorrect permission flag. The most common problems are about `/tmp` which should be 1777 and `/dev/ttys?` which might be 766. So as root do

```
chmod 1777 /tmp
chmod 766 /dev/ttys?
```

6.2.6 du reports twice the size showed with ls -l, is it a bug?

No it is not, the report is 512 bytes multiple (due to POSIX requirement), for KB you just add the `-k` flag. You can add a `du` function in your `.profile` which does this automatically, something like `du(){ /usr/bin/du -k $* }`.

6.2.7 Sometimes, I get mount: can't open lock file, what does this means?

This can happen for two reasons:

1. You try to mount something as non-root. In that case you can either retry as root, or set the setuid bit to mount as follows:

- be sure that `mount` belongs to root, if not do `chown root /bin/mount`
- set the setuid bit with `chmod u+s /bin/mount`

BTW you have to do the same with `umount` (in order to be able to unmount) Remark that it is NOT safe to allow anyone to perform `mount/umount`.

2. You are root. `mount` wants to open `/etc/mtab` and `/etc/mtab~` — the first one for reading, the second as lock file. If there is already a `mtab~` remove it. This can happen if you used once `gnuemacs` on `mtab`.

To forbid that case, just add the following entry in your `/etc/rc` file:

```
rm -f /etc/mtab~
```

6.2.8 When I try mount /dev/hd?? /user, I get error 2.

Be sure, that your mount point `/user` does exist; if not perform a `mkdir /user`.

BTW the error numbering is explained in `/usr/include/errno.h`

6.2.9 Since I have upgrade my Linux kernel, ps won't work anymore, why?

Each time you upgrade (or re-compile the kernel), you should perform a `ps -U` (`-U` is for update the `/etc/psdatabase`).

Every time you boot a new kernel you have to do a `ps -U` to update the `psdatabase`, after doing this you can remove the system file or even have to patch the `ps-suite`.

BTW: sometimes a patch makes recompiling `ps` necessary. Sometimes you do a make clean.

NOTE: As of 0.99 there is now a `ps` package which uses the `/proc filesystem`. This is a meta-filesystem which has an entry for each process that's running... create the directory `/proc` and do a `mount -t proc /proc /proc` to see it, or else add the following entry in your `/etc/fstab`:

```
none /proc proc
```

You can get the `procps` package from `tsx-11.mit.edu` in `/pub/linux/BETA/procps`. The nice thing about this is: you don't have to upgrade it for every kernel version (and thus `ps -U` is meaningless if you use `procps`).

MKJ wrote in c.o.l.a:

```
tsx-11.mit.edu: /pub/linux/BETA/procps/procps.tar.Z is the latest version of procps. Most of the reported bugs (i.e. all I can remember at the moment ;- ) have been fixed, a man page for ps has been added, and several new programs have been added and/or fixed. A list of the stuff provided:
```

```
ps
free (a real one, from Linus)
uptime (this needs work to be good, mine is a hack)
tload (a hack, but who cares?)
xload (from drew)
fuser (perl script prototype, hopefully the real thing in a later release)
```

6.2.10 Since Linux 0.96b I have a lot of core files all over my disk. How can I sweep them away?

Use the following command

```
find / -name core -exec rm {} \;
```

BTW: think twice before using this command, there might be a John Core user on your system; this command will erase his mail :)

If you never want to see another core file, put this line into your `/etc/rc` file:

```
for a bash user
    ulimit -c 0
for a csh or tcsh user
    limit coredumpsize 0
```

6.2.11 I can only log-in as root, is it normal?

No, add `"rm -f /etc/nologin"` in your `/etc/rc.local` file

6.2.12 Sometimes my Linux virtual console get messed up when I cat a binary file. The characters shown are above ascii 127, any clue?

J. Stuckey (`stuckey@mrcnext.cso.uiuc.edu`) says:

It's a vt-100 thing I think.

```
echo "^V^O"
```

should fix it. Control-V is the literal-next character that `stty` reports for me, and Control-O is the *magic* character that returns you to livability.

You can also use

```
echo "^V<ESC>c"
```

which also clears the screen.

7 Linux and DOS

This section is co-written by Mark Komarinski, A. V. Le Blanc and MM. Corsini. The official maintainer is Mark, if you have *any* questions, critics . . . , mail him at `komarimf@craft.camp.clarkson.edu`.
Last Update March 1993.

7.1 General

7.1.1 Is it possible to access DOS from Linux?

Yes.

1. The MTools package allows you to access DOS files; it emulates the DOS commands CD, COPY, DEL, DIR, TYPE, and others.
2. Since approximately version 0.97 of the kernel, you can mount DOS file systems as part of your Linux directory tree, pro
3. A DOS emulator is in alpha test, which will allow some DOS programs and utilities to run under Linux.

7.2 M-Tools

7.2.1 Why use mtools if you can just mount a drive?

Mtools is good if you want to do something quickly. For example, if you want to get directories on a bunch of floppies. The mount procedure requires you to mount the drive, get a directory, then umount it. Mtools lets you get the directory with one command.

Mtools is really fast when copying disks. I mount the SLS directory of my SUN to my Linux box and use `mcopy` to get the files on the disks. 3 times faster than using `xcopy` under DOS

7.2.2 How do I get the mtools package set up correctly?

The mtools package is available in source form on most Linux ftp sites. The most recent version (As of Sept 1992) is `mtools.n2`, and there are linked binaries and library (.a) files available as well (for example, as part of MCC interim releases of Linux). In the n2 release of mtools, there is only one executable binary which works differently depending on its name: you can create hard or symbolic links to it named `mcopy`, `mdel`, `mdir`, `mtype`, etc; this is how the Makefile in version n2 of mtools does it, and it saves several hundred kilobytes of disk space. Finally, you need the correct information in the file `/etc/mtools`, which is described below. Mtools also comes with the SLS release of Linux.

7.2.3 What is the format of the file `/verb/etc/mtools`?

A complete entry in the file `/etc/mtools` contains the following fields: drive, device, fat, cylinders, heads, sectors, offset. Two examples of entries from `/etc/mtools` are

```
A /dev/fd0 12 80 2 15
C /dev/hda1 16 0 0 0
```

which defines the DOS disk A: as accessible through the device `/dev/fd0`, having a 12-bit FAT, 80 cylinders, 2 heads, and 15 sectors per track; DOS disk C: is accessible through the device `/dev/hda1`, has a 16-bit FAT, and its geometry is

simply that of the hard disk where it lives. The last three numbers can be 0 if you wish; this allows mtools to try to figure out the disk's geometry itself, and perhaps to fail. A 12-bit FAT is common for floppies, but may occur in small hard disk partitions. A 16-bit FAT is common for hard disks. This is an extract of my `/etc/mtools` file:

```
A /dev/fd0 12 0 0 0 # 3.5 1.4 Meg (autodetect)
B /dev/fd1 12 0 0 0 # 5.25 1.2 Meg (autodetect)
C /dev/hda1 16 0 0 0 # 1st partition of my Disk
```

7.2.4 Where can I find out more about mtools?

There are two detailed `README` files in the `mtools.n2` distribution. These files treat compiling and using mtools. There is a file `README.mtools` which treats only using mtools, which is a part of the MCC interim version of Linux.

7.2.5 This sounds me like a chicken and eggs problem, how can I install the mtools package at the very beginning.

You have to use the rawrite stuff or the mount procedure.

7.3 DOS filesystem

7.3.1 How do I use the DOS file system?

The DOS file system is part of the kernel. If you have a kernel of level 97.1 or above, and an appropriate mount command, type

```
mount -t msdos [-o conv=text] /dev/hda1 /dos
```

to mount the partition `/dev/hda1` as an MS-DOS file system on the directory `/dos`. You'll need a recent mount command, from at least release 97 or later of the root disk. Recent mount commands also accept the options `conv=binary|text|auto` (default is binary) to specify that text end-of-lines in DOS files are to be converted to UNIX end-of-lines (by omitting carriage return characters) in no cases (binary) or in all cases (text) or in cases that don't have well-known binary extensions (e.g., `.EXE` or `.COM`) (auto).

7.3.2 When I mount an MS-DOS disk, I get two lines of what seem to be error messages and the message: No bmap support. Is my disk bad?

No. Since the `dosfs` is still in ALPHA mode, the two lines are used for displaying information about the disk that was mounted, just in case you have problems. The `No bmap support` message means that you cannot run Linux programs from that disk. This has been fixed in the latest release (alpha.9)

7.3.3 I want to use the DOS file system with either `conv=binary` or `conv=auto`, but I want to convert text files from DOS to UNIX format, or from UNIX to DOS format.

Use the utility `todos/ fromdos` which comes as part of the `dosfs` package, currently (Sept 92) in released in version 8, or use the `flip` utility by Rahul Dhesi.

7.3.4 Where can I find out more about the DOS file system

There is a **README** file included in **dosfs.XX.tar.Z** (the current value of **XX** is 10), and an abbreviated version of this in the MCC interim package.

7.4 DOS-Emulator

7.4.1 What about the dos emulator (dosemu)?

Dosemu is in alpha testing now, so it is very unstable and crashes easily. Some programs (such as earlier versions of Turbo Pascal; TP 5.5 too) seem to work fine, but other programs such as **dir /p** will crash the emulator.

Look in your favorite FTP site for the dosemu kit. For a semi-complete list of programs that do and do not work under the emulator, get the file **/pub/Linux/docs/emusucc.txt** at **sunsite.unc.edu** or **tsx-11.mit.edu**.

7.4.2 I have a Stacker volume on my hard drive. Can I access that?

Yes, but you will have to do that through the dosemu program. For more information on it, check out the file **/pub/Linux/docs/stacker.doc** at **sunsite.unc.edu**

7.5 Other file systems

7.5.1 How about OS/2 HPFS partitions, or MacIntosh/Amiga floppies? Can I mount any of them?

There is no support for any of those filesystems yet.

Some guys are working on the **amigafts**, the Amiga FLoppy File System. Look out for this file.

7.6 Additions

7.6.1 Could someone explain how to use rawrite?

Well, rawrite is a DOS util, which write sequential sector of a formatted disk/floppy. When a floppy has been rawritten, you can (under Linux) untar it (use **x**, **v**, **z** and **f** flags). As an example:

1. under DOS use **rawrite**

```
rawrite
source: kermit.z
destination: a
```

2. boot under Linux, and perform a tar

```
tar zxvf /dev/fd0
tar zxvp < /dev/fd0
```

You DO NOT NEED TO MOUNT a rawritten disk

7.6.2 What is as86.tar.Z?

It's the port of Bruce Evans' minix assembler, you need it to be able to recompile Linux at your convenience. In fact this is ONLY used for **boot/setup.S** and **boot/bootsect.S** they create 80x86 REAL mode code.

7.6.3 Turbo (Microsoft) Assembler won't compile the Linux boot code. In fact, some of the opcodes in these files look completely unfamiliar. Why?

The Linux boot codes are written in Bruce Evans' Minix assembler, which has the same opcodes as the original Minix assembler ported to Linux get `as86.tar.Z`. Anyway there are a few differences between these and normal DOS assemblers.

8 Miscellaneous Hints

This section is maintained by Matt Welsh (`mdw@tc.cornell.edu`). Please mail me if you have any changes/updates/questions. Thanks -mdw

Special gcc information are located in section 12. A special section is devoted to it since it's *the* compiler of Linux. I have subsectionned this part in 3 subsections: Misc Information, Major/Minor Information and Serial-Information.

8.1 Misc Information

8.1.1 It seems that \$#! ported on Linux don't run correctly, what do I do about reporting bugs?

It's possible that either the program itself has a bug or that Linux has a problem that this program brings out. :) But first check that the size of the file(s) corresponds that of the files(s) on the FTP sites where it's available. If they're different, either you downloaded them incorrectly (i.e. you forgot to turn on `bin`) or whoever put them on the FTP site uploaded them incorrectly.

If that's not the problem, then post to `comp.os.linux` asking about the program, to verify that it is a bug. PLEASE: when posting possible *bug reports* include all error and output information from running/compiling the program. Just saying *it doesn't work* isn't very helpful. Also mention your specific setup, Linux version, GCC version, etc. Some of these things depend on running under certain versions and you may have missed that information.

Note that my `ml-linux-bugs@dg-rtp.dg.com` bug reporting list has been phased out. It turns out that Linux has so few bugs, most of which are resolved on the newsgroup or through Linus before I can accumulate them and post. :) In short: if there's a bug in Linux or in Linux-ported software, it will usually be fixed in the next patchlevel or version.

8.1.2 Has \$#! been ported to Linux?

First check out the FTP sites and read the monthly INFO-SHEET, as well as the new *Linux News* and the META-FAQ's (all of which are either available on the FTP sites and/or posted to the newsgroup as they're written). Also check out the *Linux Project Registry* (posted to the newsgroup and on the FTP sites) which lists ongoing/current Linux projects. Also look in the *old* Linux digests and mailing-list archives, kept on `tsx-11.mit.edu` and `nic.funet.fi`. Also, see if there's a GNU³ version of the program you're looking for (which are available everywhere). Since Linux uses GCC as its native compiler, most GNU software ports directly to Linux without problems. If all else fails, ask on the mailing list or newsgroup if the program is ported and where it's available.

³GNU stands for GNU's Not Unix, which (besides being a recursive acronym) is a project started by the Free Software Foundation (the FSF) to write a freely distributable version of UNIX. The GNU kernel is named HURD, and is based on Mach. It is currently being written, and is not yet done. Many of the GNU utilities, however, are completed and are much more functional than the original Unix utilities. Since they are freely available, Linux is using them as well.

8.1.3 I've ported \$#! to Linux, what should I do to add it in the standard distribution?

First read the previous section, then talk to the maintainers of the various releases (MCC interim, SLS, and so on) about including your program with those releases. The best way to make programs available to the rest of the world is to upload it to one of the Linux FTP sites (such as `tsx-11.mit.edu: /pub/incoming` or `sunsite.unc.edu: /pub/Linux/Incoming`).

BTW Whenever you submit binaries, please think that if you link them with shared libraries, this might cause problems for those who (gasp!) don't have the shared libraries installed. You can either link them using `-static`, and if someone wants to build a shared version on their own machine they can get the sources (which you should also make available) and build it themselves. Or else provide the shared lib.

8.1.4 I want to port \$#! to Linux, what are the flags?

Recall that Linux implements subset of SYSV and POSIX, so `-DUSG` and `-DPOSIX` work in general.

NOTE1: `SIGBUS` is not there, and can be safely commented out in general.

NOTE2: see section 12 for more details.

8.2 Minor/Major Information

This section is maintained by Rick Miller, any comments, suggestions, remarks should be mailed to him at `rick@ee.uwm.edu`.

Last Update April 1993.

8.2.1 What are the Major/Minor device numbers

This is a list of the device names, along with Major and minor numbers, which the Linux kernel may currently recognize. I say *may* because some of them may require patches to your kernel, others must be configured-in using the Linux kernel's `make config` routine before compiling.

My profuse thanks to all the writers of device drivers for Linux, and to all those patient folks who took the time to explain their devices to me. May they become obscenely rich. Until then, I hope this list will help prevent driver collisions and perhaps lend a clue to a newbie or two.

IF YOU WANT ANY CHANGES OR ADDITIONS TO THIS LIST, TELL ME!

DEVICES NOT LISTED HERE SHOULD USE MAJOR NUMBERS ABOVE 127 UNTIL ALLOCATED A MORE PERMANENT NUMBER IN THE LOWER RANGE. (I suggest using a more-or-less random number to avoid the chance of collisions with any other experimental drivers.)

TO HAVE MAJOR/MINOR NUMBERS ALLOCATED (OFFICIALLY) FOR YOUR DEVICE DRIVER, SEND E-MAIL TO: `rick@ee.uwm.edu`

Major Devicees

Value	Name	Type	Used for
0	Unnamed	(unknown)	for proc-fs, NFS clients, etc
1	Memory	(character)	ram, mem, kmem, null, port, zero, core
2	Floppy	(block)	fd[0-1]<[dhDH]{360,720,1200,1440}>
3	AT-Disk	(block)	hd[a-d]<[0-8]>
4	Tty	(character)	tty's and pty's
5	Dial-out	(character)	tty, cua[0-63]
6	Parallel	(character)	par[0-2], lp[0-2]
7	Unused		
8	SCSI-Disk	(block)	sd[a-h]<[0-8]>
9	SCSI-Tape	(character)	<n>st[0-1] or <n>rmt[0-1]
10	Mouse	(character)	logibm, psaux, inportbm, atibm, (mouse)
11	CD-ROM	(block)	scd[0-1]
12	QIC-Tape?	(character)	rmt{8,16}, tape<{-d,-reset}>
13	XT-Disk	(block)	xd[a-b]<[0-8]>
14	Audio	(character)	audio, dsp, midi, mixer, sequencer
15	Joystick	(?)	js0, js1
16	Socket	(character)	net, arp
17	AF_UNIX	(character)	unix
18	AF_INET	(character)	inet, ip, icmp, tcp, udp
19	WE-driver	(character)	we[0-3]
20	DP8390-driver	(character)	wd[0-3], ec[0-3], ne[0-3]
21	Sony CD-ROM	(block)	sonycd
22	2nd IDE-Adapter	(block)	hd1[a-b]<0-8>
23	Mitsumi CD-ROM	(block)	mcd
31	Link Interface	(character?)	??

Breakdown of minors by Majors

0. Unnamed Minors?

1. Memory

0	/dev/ram	a BLOCK device (RAM-disk)
1	/dev/mem	
2	/dev/kmem	
3	/dev/null	
4	/dev/port	
5	/dev/zero	
6	/dev/core	like /dev/mem, but in "core"-file format for gdb

2. Floppy Minors are [drive + [4 * type]] where drive 0-3 == A:-D: (floppy) and type is:

0	Autodetect
1	360k on 5.25" DD
2	1.2M on 5.25" HD
3	360k on 3.5" DD
4	720k on 3.5" DD
5	360k on 5.25" HD
6	720k on 5.25" HD
7	1.44M on 3.5" HD

Value	Device	Usage
0	/dev/fd0	Autodetected first floppy
1	/dev/fd1	Autodetected second floppy
2	/dev/fd2	Autodetected third floppy
3	/dev/fd3	Autodetected fourth floppy
4	/dev/fd0d360	360k on 5.25" DD in first drive
5	/dev/fd1d360	360k on 5.25" DD in second drive
6	/dev/fd2d360	360k on 5.25" DD in third drive
7	/dev/fd3d360	360k on 5.25" DD in fourth drive
8	/dev/fd0h1200	1.2M on 5.25" HD in first drive
9	/dev/fd1h1200	1.2M on 5.25" HD in second drive
10	/dev/fd2h1200	1.2M on 5.25" HD in third drive
11	/dev/fd3h1200	1.2M on 5.25" HD in fourth drive
12	/dev/fd0D360	360k on 3.5" DD in first drive
	/dev/fd0H360	- " -
13	/dev/fd1D360	360k on 3.5" DD in second drive
	/dev/fd1H360	- " -
14	/dev/fd2D360	360k on 3.5" DD in third drive
	/dev/fd2H360	- " -
15	/dev/fd3D360	360k on 3.5" DD in fourth drive
	/dev/fd3H360	- " -
16	/dev/fd0D720	720k on 3.5" DD in first drive
	/dev/fd0H720	- " -
17	/dev/fd1D720	720k on 3.5" DD in second drive
	/dev/fd1H720	- " -
18	/dev/fd2D720	720k on 3.5" DD in third drive
	/dev/fd2H720	- " -
19	/dev/fd3D720	720k on 3.5" DD in fourth drive
	/dev/fd3H720	- " -
20	/dev/fd0h360	360k on 5.25" HD in first drive
21	/dev/fd1h360	360k on 5.25" HD in second drive
22	/dev/fd2h360	360k on 5.25" HD in third drive
23	/dev/fd3h360	360k on 5.25" HD in fourth drive
24	/dev/fd0h720	720k on 5.25" HD in first drive
25	/dev/fd1h720	720k on 5.25" HD in second drive
26	/dev/fd2h720	720k on 5.25" HD in third drive
27	/dev/fd3h720	720k on 5.25" HD in fourth drive
28	/dev/fd0H1440	1.44M on 3.5" HD in first drive
29	/dev/fd1H1440	1.44M on 3.5" HD in second drive
30	/dev/fd2H1440	1.44M on 3.5" HD in third drive
31	/dev/fd3H1440	1.44M on 3.5" HD in fourth drive

Naming goes like this:

fd[drive][media][size]

where: [drive]=0-3

[media]={d,h,D,H}:

Corresponds to DOS's "A:"-"D:"

d=Double Density 5.25" diskette

h=High Density 5.25" diskette

D=Double Density 3.5" diskette

H=High Density 3.5" diskette

[size]={360,720,1200,1440} kilobytes

Floppies are assumed to be double-sided (DS), and drives are assumed to be high-density devices.

3. AT-Disk (For IDE, MFM, and RLE drives and controllers.)

- On the first AT controller card

0	/dev/hda	The whole first HD, including its MBR
	/dev/hda0	- " -
1-4	/dev/hda\{1-4\}	Primary partitions on the first hard drive
5-8	/dev/hda\{5-8\}	Extended partitions on the first hard drive
64	/dev/hdb	The whole second HD, including its MBR
	/dev/hdb0	- " -
65-68	/dev/hdb\{1-4\}	Primary partitions on the second hard drive
69-72	/dev/hdb\{5-8\}	Extended partitions on the second hard drive
- On the second AT controller card

128	/dev/hdc	The whole third HD, including its MBR
	/dev/hdc0	- " -
129-132	/dev/hdc\{1-4\}	Primary partitions on the third hard drive
133-136	/dev/hdc\{5-8\}	Extended partitions on the third hard drive
192	/dev/hdd	The whole fourth HD, including its MBR
	/dev/hdd0	- " -
193-196	/dev/hdd\{1-4\}	Primary partitions on the fourth hard drive
197-200	/dev/hdd\{5-8\}	Extended partitions on the fourth hard drive

Notes: BE **VERY CAREFUL WITH** the four *whole drive* devices (`hda`, `hdb`, `hdc`, and `hdd`)!! These four devices embody the *entire drive*, not just one partition. The only things that use them are things that need to read/change the partition table (like `fdisk`).

Linux doesn't order anything. It perceives partitions in the order in which they appear in the partition table. Thus, `/dev/hd?1` may follow `/dev/hd?2` in the cylinder numbering.

The names of the hard drives are not the same as under Minix.

4. Tty

0	/dev/tty0	This is the currently active Virtual Console
1-63	/dev/tty[1-63]	Specific virtual consoles
64-127	/dev/ttyS[0-63]	Serial ports (dial-in mode)
128-191	/dev/pty[p-s][0-f]	PTY Masters
192-255	/dev/tty[p-s][0-f]	PTY Slaves. ([0-f]=0123456789abcdef)

Notes: Notes: NOTICE THE NEW NAMING FOR SERIAL LINES. Serial lines will be named either `tty...` or `cua...` (See Major #5) followed by the one-digit number of the *board* the line is from, and the number of the line on that board. The four default lines are considered to be on board "0", so what DOS calls "COM2:" should now be known as `/dev/tty01` or `/dev/cua01`.

For more information on the use of multi-port boards, read the documentation which accompanies the serial patches.

There are several constants set in the kernel sources which can be changed to compile a more customized kernel. They're found in (`/usr/src/linux/include/linux/tty.h`):

`NR_CONSOLES` The number of virtual consoles

`NR_SERIALS` The number of serial lines

5. Dial-out

0	/dev/tty	the tty that owns the process calling it
64-127	/dev/cua[0-63]	Serial ports (dial-out mode)

See Major 4 for new naming convention.

6. Parallel Parallel (printer) ports. (Increaseable in `include/linux/lp.h`)

```
0 | /dev/lp0
1 | /dev/lp1  Same as MS-DOS's LPT1: on my machine...
2 | /dev/lp2
```

Notes: The number of line printers is defined by `LP_NO` which is found in `[usr/src]/linux/include/lp.h`.

Even if you only have one printer port, Linux may still call it `/dev/lp1` (rather than `/dev/lp0`). Watch during boot-up to see how it recognizes your printer port.

7. Unused (first come, first serve) This one's getting old. No minor numbers are yet assigned. It's not even in the source code. Maybe it never will be...

8. SCSI-Disk Minors numbers are $((16 * Drive) + Partition)$ where *Drive* is the number of the physical drive in order of detection and *Partition* is as follows:

```
0 | is the whole drive
1-4 | are the DOS primary partitions
5-15 | are the DOS extended (or logical) partitions, so...

0 | /dev/sda      The first (detected) SCSI drive
   | /dev/sda0    - " -
1-4 | /dev/sda[1-4] Primary partitions on the first SCSI drive
5-15 | /dev/sda[5-15] Extended partitions on the first SCSI drive
16 | /dev/sdb     The second (detected) SCSI drive
   | /dev/sdb0    - " -
17-20 | /dev/sdb[1-4] Primary partitions on the second drive
21-31 | /dev/sdb[5-15] Extended partitions on the second drive
32 | /dev/sdc     The third (detected) SCSI drive
   | /dev/sdc0    - " -
...and so on
```

9. SCSI-Tape

```
0 | /dev/st0  First (detected) SCSI tape drive, rewind-on-close
1 | /dev/st1  Second (detected) SCSI tape drive, rewind-on-close
128 | /dev/nst0 First (detected) SCSI tape, no rewind-on-close
129 | /dev/nst1 Second (detected) SCSI tape, no rewind-on-close
```

10. Mouse (MGR may require that `/dev/mouse` be linked to one of these...) NOTE THE CHANGES TO DIFFERENTIATE MOUSE TYPES! (Please implement these in the distributions.)

```
0 | /dev/logibm  Logitech-'compatible' bus mouse
1 | /dev/psaux  PS/2 mouse port (may not work on some lap-tops, yet)
2 | /dev/inportbm  Microsoft "InPort" bus mouse
3 | /dev/atibm   ATI XL bus mouse
```

11. CD-ROM

```
0 | /dev/scd0  The first (detected) SCSI CD-ROM
1 | /dev/scd1  The second (detected) SCSI CD-ROM
```

(*There's not much more to it than that says Eric Youngdale.*)

12. QIC-Tape? (I really don't have much info on this one...)

6		/dev/rmt8	QIC-120
8		/dev/rmt16	QIC-150
		/dev/tape	- " -
136		/dev/tape-d	(It has something to do with being 128+8... ?)
255		/dev/tape-reset	for resetting only

13. XT-disk XT (8-bit) hard disk controller devices.

Minor numbers are assigned in the same manner as for the normal (AT-type) Hard Drive devices (/dev/hd*).

14. Audio

0		/dev/mixer	Mixer and Control Device
1		/dev/sequencer	FM-synthesizer and Midi
2		/dev/midi	(for future use)
3		/dev/dsp	Digitized voice (DAC/ADC)
4		/dev/audio	(Reserved for compatibility with Sun)

15. Joystick

0		/dev/js0	(Left/Right?) joystick
1		/dev/js1	(Right/Left?) joystick

16. Socket

0		/dev/net	Generic layer (sockets)
1		/dev/arp	Address Resolution Protocol (ARP)

17. AF_UNIX

0		/dev/unix	AF_UNIX protocol layer
---	--	-----------	------------------------

18. AF_INET

0		/dev/inet	AF_INET protocol layer
1		/dev/ip	Internet Protocol (IP)
2		/dev/icmp	Internet Communications Management(?) Protocol (ICMP)
3		/dev/tcp	Transmission Control Protocol (TCP)
4		/dev/udp	Protocol (UDP)

19. WE-Driver

0-3		/dev/we[0-3]	"WE" driver
-----	--	--------------	-------------

20. DP8390-Driver

0-3		/dev/wd[0-3]	"DP8390" driver, WD8003
8-11		/dev/ec[0-3]	"DP8390" driver, 3C503
16-19		/dev/ne[0-3]	"DP8390" driver, NE2000

21. Sony-CD-ROM

0		/dev/sonycd	Sony CDU-535 CD-ROM, I think.
---	--	-------------	-------------------------------

22. 2nd AT-Cntrlr

Minor numbers are like those of the primary AT-Disk controller. THIS MAY WORK FOR HARD-CARDS. (Note the "1" in the device name for the secondary controller. This is because we may eventually wish to change the names of the primary-controller drives to "hd0[a-b]...".)

23. Mitsumi CD-ROM

I don't know much about this one.

31. Link Interface

Link Interface for use with Christoph Niemann's driver for the INMOS C012-based high-speed serial link interface card (useful for talking to transputers or video digitizers).

NOTE: all the numbers given are in decimal form (the one you can see if you perform `ls -l` on `/dev`).

8.2.2 Could someone give me indication about the meaning of the IRQ's?

This answer comes from `alawrenc@sobeco.com`.

Standard IRQ's:

IRQ 0	Timer
IRQ 1	Keyboard
IRQ 2	Slave 8259 (AT)
IRQ 3	COM 2 / COM 4
IRQ 4	COM 1 / COM 3
IRQ 5	(XT) Hard Disk, (AT) LPT2
IRQ 6	Floppy Disk
IRQ 7	LPT1
IRQ 8	Real Time Clock (AT only)
IRQ 9	Re-direct to IRQ 2
IRQ 10	Not Defined
IRQ 11	Not Defined
IRQ 12	Not Defined
IRQ 13	80287
IRQ 14	Hard Disk (AT)
IRQ 15	Not Defined

8.2.3 Could some one clear up the devices meaning?

`port` This allows programs to access the hardware ports directly. Not something you generally mess around with much.

`ptyp0-3ptypa...tty` These are the pseudo-tty *master* devices. Each pty connection uses a slave-master set of tty devices.

`tty0...` `tty[1-8]` are the virtual consoles associated with Alt-F[1-8]. `tty0` is the current virtual console (so writing something to `tty0` goes to the current vc).

`tty64` I've figured out is the modem connection

Yes, that would correspond to COM1 under DOS. However, the `tty64` name is obsolete - `ttys[1-4]` should be used instead.

`ttyp0...;ttypa...` These are the pseudo-tty *slave* devices.

`ttys1...` These are the serial devices. `ttys1` corresponds to COM1 under DOS, `ttys2` corresponds to COM2, etc.

8.3 Special Serial

This section is maintained by Jim Gifford (`jgifford@{attmail,world.std}.com`), some addenda were sent by Juha Ursin (`jus@snakemail.hut.fi`).

Last update May 1993.

Douglas E. Quale:

This has come up a couple of times already (including the case of serial mice as well), but for the record stty acts on stdin not stdout. Old stty's (from V7 through BSD4.3) used stdout, but this is suboptimal and doesn't conform to POSIX. The GNU stty you are likely using on Linux uses stdin, as does the stty distributed with BSD Networking Release 2. (Also, `stty -a` is more informative about possible parameters, although it's pretty hard for me to remember what 90

Jim Gifford:

However, there are a few older (of mysterious origin) stty's that work on stdout (I have one myself!)

8.3.1 Is there a list somewhere where I can get help with serial communications under Linux?

There is a list for the discussion of serial communications under Linux. It is for problems, drivers, new developments, etc... with the Linux serial devices. The list is: `linux-serial@stolaf.edu` To join, send mail to `linux-serial-request@stolaf.edu` I hope that this list will prove beneficial to the improvement of Linux. This list is maintained by Michael K. Johnson as `linux-serial-request@stolaf.edu`.

8.3.2 When I run kermit under Linux, I get "Warning, Read access to lock directory denied". What am I doing wrong?

Nothing, you just need to create `/usr/spool/uucp` (kermit 4.6?) or `/usr/spool/locks` (this is for the kermit5A), which is where kermit like to lock files.

8.3.3 What are the major, minor numbers for the serial ports under linux?

```
64 /dev/ttyS0 com1
65 /dev/ttyS1 com2
66 /dev/ttyS2 com3
67 /dev/ttyS3 com4
```

8.3.4 Can anyone give me a sample /etc/inittab entry for login from a pc attached to serial line /dev/ttyS2?

Humberto speaking :), updated by Rick Miller:

First set up the modem to turn off echo and enable auto answer, I do this in kermit by connecting to the modem and typing `ate0s0=1` followed by enter (w/o quotes). Then setup inittab to spawn getty on the modem

```
ttys2:vt100:/etc/getty -m 1200 ttys2
```

(Replace `vt100` with the name of the `/etc/termcap` entry for the terminal type you will use, or use `dumb` if you don't have one.)

Then it should work. Some modems can be permanently set to disable echo and set auto answer, see your manual.

Jim Weigand says:

disable all messages. This will prevent getty from hanging up your modem.

Set	For
ATE0	No echo
ATQ1	No messages
ATS0=2	Answer 2nd ring
ATS7=60	1 minute to answer (shorter if 2400 baud)

You can use `kermit` to set these. Do an **AT&W** to save for power-up.

Michael K. Johnson says:

If you would rather not save these commands as defaults to come up on power-up, perhaps because you want to use your current modem settings under a DOS communications package, you can also shove these command out `ttys?` from `/etc/rc` (or `/etc/rc.local`) using the command:

```
echo "<modem_settings>" > /dev/ttys?
```

8.3.5 How do I set parameters like parity for serial login?

Use `stty` and redirect input from the serial line.

```
ex: stty parodd < /dev/ttys2
```

which gives `ttys2` odd parity. type `stty -a` to get an idea of possible parameters. Or `man stty` as well :)

8.3.6 I run SLS pl6, kernel is pl8 and I use libc.4.3.2. Trying to allow login on my terminal (on `ttyS1` with only three wires (RX, TX, GND) connected). `Getty 9600 /dev/ttyS1` starts and I get the login prompt, but the password prompt never appears.

Ted Ts'o:

Setting `CLOCAL` is the correct way to solve this — it is not a bug in the kernel. Since you only have three wires connected, it means that the carrier detect line was left floating, and probably floated to ground. Thus, the `/dev/ttyS1` device blocked waiting for carrier detect to go high. The reason why the login prompt appeared is that `getty` opened `/dev/ttyS1` with the `NODELAY` flag; but when `/dev/ttyS1` was reopened for `/bin/login`, it did not have the `NODELAY` flag, and you ended up blocking.

That's what the `CLOCAL` flag is for — it indicates that the serial port is hooked up to a local (hence `CLOCAL`) terminal, and so the carrier detect line should be ignored. In general, whenever the carrier detect line is not supported by a serial device, or if your RS-232 cable does not support the CD wire, the `CLOCAL` flag should be set.

8.3.7 Where are my Call-out-mode-serial ports (`cuaX`)?

When should I use these `cuaX` ports ?, Why on my system `ttys3` is used instead of `ttys3?`

Ask the person who created the install procedure (Peter MacDonald, in the case of SLS).

When should I use these `cuaX` ports?

When you're dialing out (i.e., `kermit`, `cu`, `seyon`, `uucp`, etc.).

On my configuration `ttys3` is named `ttys3` incorrectly — why?

Again, see the person who created the install procedure. Recent versions of SLS create the serial devices using the new convention (`ttys3`). `ttys3` represents an older convention, and some older install procedures may not have switched over.

8.3.8 What other possibly free IRQ's there are than `IRQ5`?


`IRQ 2` is also often free, although some networking cares use them.

8.3.9 Can I use IRQ 7 when not printing?

It's generally a bad idea to have two cards plugged in sharing an IRQ address. It happens to work with most serial cards, but it doesn't always work. You can use IRQ 7 if you don't have a parallel port plugged into your system. If you do have a parallel port plugged in, you can try using IRQ 7; it might work, it might not. Be sure to also make sure that the printer works after you plug in your serial card configured to use IRQ 7. Another alternative: If you are using the version of the parallel printer driver that doesn't use interrupts, you can go to your parallel port board, and disable the IRQ, thus freeing IRQ 7.

8.3.10 Is pin 1 nearest to the plate or is it 31?

If this is a sample PC ISA card, then A1 is the pin closest the *back* of the card, and on the top, if the pins are held down and the *back* of the card is to your right:



The "B1..B31" pins are on the reverse side of the card.

9 More Hints

This part try to keep track of the different information that appeared in `comp.os.linux` and on the list since beginning of March. I tried to update it for 0.99, so there might be some mistakes. Moreover take care to use the correct library and include stuff, and the ad-hoc gcc you use !!!

9.1 Booting

9.1.1 How to start Linux from drive B?

There is a DOS utility called `boot_b.exe` (look at DOS ftp). Another simple way is to open the box and invert the cables.

9.1.2 The program `boot_b` works fine *but* once the first disk is read the system goes back to the first drive, any hints?

Yes, change the bootimage in just the same way that you change it to boot on the hard drive, except that the major/minor pair is different. All these information are in the file `INSTALL-0.10`. Remember that if you use a sun or other endian machine, you will need to reverse the byte order when you run the filter program (also in the same file).

9.1.3 How can I get Linux to boot directly from the harddisk?

The best option right now is LILO version ALPHA.9. It has been generally agreed that the days of using `shoelace` are at an end. See section 5.3.8 of this FAQ for LILO information.

9.1.4 I use shoelace but I want to change my root partition, what is the process to get rid of it?

With Norton utility you can put back a standard boot sector. Another possibility is to restore the old boot sector (the one you should have backup *before* installing `shoelace`).

9.1.5 Is it possible to disable the 3-fingers salute (ctrl-alt-del)?

Yes, in `kernel/sys.c` you can read the following:

```
/*
 * this indicates wether you can reboot with ctrl-alt-del: the deault is yes
 */
static int C_A_D = 1;
```

there is also a small utility written by Linus in `digest242 vol#2`

9.1.6 Could some one explain the information provided at boot-time?

```
serial port at 0x03f8 is a 16450
serial port at 0x02f8 is a 16450 (what's that the uart chip?)
```

Right, the last number should either be 8250, 16450, 16550, or 16550a, and on the two 16550 models, it will report that FIFO's have been disable (16550) or enabled (16550a).

```
8 virtual consoles (that's how many alt-F's I can get going?
but only F1-4 actually work)
```

You can get sessions running on Alt-F[1-8], but the *standard* `/etc/inittab` only runs `getty/login` on Alt-F[1-4]. You can start sessions on the other consoles by using `doshell` or adding lines to `/etc/inittab`.

```
4 pty's (are these the consoles F1-4?)
```

No, those are *pseudo* ttys, which programs like MGR use to simulate tty connections. That's probably a gross over-simplification, but it gives you the general idea, I think. :)

```
p_init: lp1 exists (0) (is that the (1)ine (p)rinter?)
```

Right.

9.1.7 What is the rdev program provided in the images?

It's a program from Werner Almesberger of ETH Zurich. With no argument it prints the first line of `/etc/mtab`. With one argument, a boot-image, it prints the device configured as the root device, and with two arguments, a boot image and a device, it sets the device as default root in the specified image.

9.2 Miscellaneous

9.2.1 How can I backup my Hd under Linux?

I know at least two ways. One possibility is `tar` and `mtools`, another possibility is the `diskbackup/ diskrestore` of Diamano Bolla (digest37 vol. #2) which saves big hd to floppies using the stdin/stdout. These utilities have been uploaded to the major sites in file `disksplit.tar.Z`.

An example usage (Roger Binns) is:

```
tar cvf - bin dev usr etc .. | compress | diskbackup
```

and to restore:

```
diskrestore | uncompress | tar xvf -
```

BTW: if you are on Ethernet you could send your files via `tar..|rsh (tar...)` or even via NFS to a host which is regularly backed up !

9.2.2 Where is which?

It depends on the sh you are running:

in bash `type -path`

in tcsh it's a builtin

for rc you can try the following (untested by me) script from Kevin Brown:

```
#!/bin/sh
for i in `echo $PATH | sed 's:/ /g'` ; do
  for j in ""$@" ; do
    if test -x "$i/$j" ; then
      echo "$i/$j"
    fi
  done
done
```

9.2.3 How to use setterm: for the novice?

he `setterm` utility provides access to most of Virtual Consoles (VCs) functionality. You can set your screen up to blank at 10 minutes using:

```
setterm -blank 10
```

You can set colors, and clear the screen. For a full list of commands, just type `setterm` with no arguments.

There are a few tricks with the screen dumper can really make VCs go a long way. Here are a few of the common ones that I use:

```
setterm dump
```

Dumps the contents of the current VC to `screen.dump` (in the current dir).

```
setterm dump 4
```

Dumps the contents of VC 4 to `screen.dump`

```
setterm -file mydumpfile -dump 4
```

Dump the contents of VC 4 to the file `mydumpfile`

```
setterm -file /dev/tty0 -dump 4
```

Dumps the contents of VC 4 to the current VC.

```
setterm -file /dev/tty4 -dump
```

Dumps the contents of the current VC to VC 4.

```
setterm -file /dev/ttys1 -dump
```

Dumps the contents of the current VC to the serial port. Handy if you are logged on and want to paste a screen full without having to resort to doing a file transfer.

```
setterm -file mydumpfile -append 4
```

Appends to instead of overwriting the dump file. Useful if you have several screens you wish to concatenate.

9.2.4 I've tried `clear/ reset` which exist on most of UNIX but it doesn't work, have I missed something?

`setterm -clear` or `setterm -reset` will solve your missing. For `clear`, you can also write a small script (which use the `cl`: part of `/etc/termcap` wrt your `TERM`), or use `bash` where `ctrl-l` will do it for you.

9.2.5 I know there are VC, but where is the `setterm` stuff?

It's in the current distribution (i.e. on the images), the source can be found in `virtcons.tar.Z` at `nic`.

9.2.6 I know there are shared libraries; does there exist an easy way to check an executable for sharing?

(Josh Yelon:) An executable which shares a library is linked with an (ordinary, non-shared) *stub* version of the library. One of the first thing this stub does (when the executable is run) is to ask the kernel to load the (big) *shared version* of the library (which is usually named `/lib/lib.XX.XX`) The upshot of this is that in the code for the stub (part of the executable), is the string `/lib/lib.XX.XX`; which can be searched by using `strings` or `grep`.

(HJ. Lu:) If you have `gcc2.11a` or later the shared image is changed to `/lib/libxxxx_vyy_zzz`. And you should better use `nm` to find *shared.lib* (`nm` failed on stripped executable). You can also write a function for `file`, which can even check the version number

(Claude Morin:) There exists at `tsx-11 ldd.tar.Z`. If you follow the instructions, you will be able to type `ldd <executable>` to List the Dynamic Dependencies of the executables. In other words `ldd` prints the name of the shared libraries needed by the executable, nothing appears for static one.

9.2.7 Sometimes, when I want to remove a directory, I get an error message, is it a (known) bug?

No, There is no bug at all, you probably have another shell on another VC whose working directory is either the one you try to remove, either a subdirectory of it.

9.2.8 I'm looking for `init`, `getty`, `login`, `passwd` stuff, where can I find them?

You should find it in `shadow.tar.Z` (only sources), at least at `tsx` in the `usr.bin` directory. Many people have reported some troubles with the **old* `shadow-passwd` (`shadow-bin.tar.Z` and `shadow-src.tar.Z`, so do not use them anymore); an alternative might be the `mcc-interim` which contains standard `passwd` binary. There is also the Peter Orbaek's `admutil-1.?.tar.Z` and `poeigl-1.?.tar.Z` which contains source for `shutdown`, `su`, `chsh`, `passwd` and a system V `init` compatible.

9.2.9 How can I setup a user account other than root?

You can either use the `adduser` program, either do it manually. In the later case, you have to:

1. edit `/etc/passwd` as root and add a line of the following format:

```
user:passwd:uid:gid:user name:home directory:login shell
```

`user` is the login name

`uid` is the numeric user id, it should be unique

`gid` is the numeric group id, this number should correspond to an entry in `/etc/group`. The `passwd` field should be left blank 'cause it is stored in an encrypted form [to set this field just use the `passwd` program].

example

```
faq::200:5:Marc-Michel:/home/faq:/bin/sh
```

2. Still as root, you should now create the home directory and set the correct ownership.

```
mkdir /home/faq
chown faq /home/faq
chgrp 5 /home/faq
```

This command is obsolete, you have to use `useradd`. To remove a user there is the command `userdel` and modifications can be made with `usermod`.

9.2.10 I've been trying to get Linux to run on my [3/4]86 box. It can't even boot. Any suggestions?

The most common error/problem is writing the bootimage to a low density disk. It fits, but the bootstrap code will only recognize high density disk. So try to format explicetely disk as high density:

- for 3.5", format `a: /n:18 /t:80`
- for 5.25", format `a: /n:15 /t:80`

9.2.11 Does there exist games, languages (other than C), and anything which make the system more friendly?

Yes, among other things there are `rogue` and `yahtzee`; `TEX`; `Prolog`, `Perl`.. but in general, if you want some extra tool port it to Linux. This is also a good beta-testing exercise.

9.2.12 Whenever I use uemacs 3.1X on a symlink, the symlink does not exist anymore, why?

Since ue3.10, uemacs uses *safe save* mode, writing the file to a temporary and moving it OVER the original. In the process, this deletes the original. To prevent this just add the following in your emacs `.rc` file: `set $ssave FALSE`.

9.2.13 Uemacs doesn't work anymore with 0.95a, whenever I want to save a file; what can I do?

`^S` and `^Q` are used for flow control. One solution is `^X^W` followed by the filename, or `M-X save-file`, try also `^XS` it could work for some version (not mine). Another possibility, if you have download the `stty.tar.Z` file, is to do `stty -IXON` before you first use uemacs (this can be included in your `.profile`). And the last is to recompile the Peter Orbaek `init-1.2` package.

9.2.14 I have an SVGA, but Linux detect an EGAc/EGAm; is it normal?

This is correct actually. You have an EGA+ card (SVGA) with a Color/Mono monitor. The only four possibilities are EGAc, EGAm, *MDA and *CGA (according to the code in `kernel/chr_drv/console.c`).

The true test, if Linux detects your video card, is if you press RETURN at the `Press <RETURN> to see SVGA- ...` boot-time message. If you have a SVGA recognized card, it will ask you to choose a screen size. If not detected, the default is 80x50 mode. BTW if you have no SVGA, press the [space] and you are in 80x25 mode.

If you have downloaded the kernel, you can automatically skip this query at boot-time if you set the `SVGA_MODE` variable in the main Makefile before compiling a new bootimage.

9.2.15 How can I change the keyboard repeat rate?

In `boot/setup.S` there are the lines:

```
! set the keyboard repeat rate to max
  mov ax,#0x0305
  mov bx,0x0000
  int 0x16
```

If you don't want to change the repeat rate at all, just comment out these lines. If you want something in the middle, change the

```
mov bx,0x000
```

```
by
```

```
mov bx,0x??yy
```

where `??yy` is determined by (Ralf Brown's interrupt list) `bh=delay value` (0x00 = 250ms to 0x03= 1000ms (one sec)) this is the delay before the repeat is performed `bl=repeat rate` (0x00 =30/sec to 0x1f=2/sec; 0x0c=10/sec [default])

9.2.16 I compiled `fdformat.c` and ran it on 1.44Mb and 1.2Mb, the results are unreadable, any clue?

`fdformat` only low-level formats the disk. to use the fdformatted disk with DOS filesystem, run `mformat` on the disk. `Mformat` writes DOS filesystem information but is unable to low-format :). In order to put a Linux filesystem on a (low)formatted disk you have to `mkfs` it.

9.2.17 What is the meaning of files ended by `.T.Z` (or `.taz`)?

The suffix `Z` is for compressed files. To uncompress them use the command

```
uncompress file.Z"
```

The suffix `T` indicates a *tar file* the usual suffix is `tar` but, the 14 chars filename limit of the Minix filesystem makes it cumbersome to use `.tar.Z` (to untar a file ended by `.T`, use the command

```
tar options file.T
```

(see the man page for more details).

For the `.taz` file, change them as `.T.Z` and go-ahead.

9.2.18 hat is the meaning of files ended by `.T.z` (or `.tpz`)?

The suffix `z` is for files compressed with the `gzip` utility, which is being used more frequently because it provides better compression. These can be uncompressed with

```
gzip -d file.z
```

Sometimes you have `zcat`, which is a link to `gzip`. Then you can also use

```
zcat file.T.z |tar xvf -
```

9.2.19 I have upgraded the kernel from `XX` to `YY` ($XX > YY$), however the login screen still says `YY`.

Just change the message in one of these files: `/etc/issue` and `/etc/motd`. The former contains the message before the login, the later is the one after you are logged in.

9.2.20 What is `doshell` good for?

It's an old program provided in the early Linux version (0.12) when the `getty` was not already there, it spawns a shell on any `tty`'s.

(Matthew Gream): I do this quite often (`getty` on `tty1`, 2, 3) and my `rc.local` spawns a shell on `tty4` as follows

```
/usr/bin/doshell /dev/tty4 /bin/tcsh &
```

9.2.21 I don't have the kernel sources, how can I change the keyboard language?

You can use the `fixkbd` program written by Laurent JULLIARD (`Laurent_JULLIARD@grenoble.hp.com`). Its purpose is to fix the keyboard map used in your kernel image. It works more or less a la `rootdev` (or `rdev`). It can be found on `tsx-11.mit.edu` in `/pub/linux/binaries/sbin/fixkbd.tar.Z`.

9.2.22 Does there exist a possibility to track down the source of kernel panic messages?

Here's info from Linus on how to track down the source of kernel panic messages. The values for EIP, address, and so on are JUST EXAMPLES: the EIP, EFLAGS, etc. can be anything.

- `Unable to handle kernel paging request at address C000000A`

This means something in the kernel tried to access kernel memory at address 0x00A (the C0000000 offset is due to the kernel starting at virtual address 0xC0000000). The first page of kernel address space is unmapped to find these kinds of incorrect references (it's probably a NULL dereference to a structure member at offset 0xA).

- `Oops: 0002`

Error code. This tells you if it's a read or write (forget which right now, and don't have any books handy)

- `EIP: 0008:000290D2`

This is where it happened: segment 0008 (kernel code segment) and offset 0x290D2 into the kernel. So to find this, you generally do something like

```
nm /usr/src/linux/tools/system | sort | less
```

and search for the function that contains the address 0x290D2.

- `EFLAGS: 00010087`

`fs: 0017`

Eflags and fs when the exception happened. You usually don't need this.

- `base: C0000000, limit: 000A0000`

`Pid: 0, process nr: 0`

In what process the exception happened. This is the swapper, which also means the thing is fatal (other processes you can just kill). That it's in the swapper process probably means that it's an interrupt that did it, as the swapper doesn't really do anything.

- `89 50 04 c7 03 00 00 00 00 c7`

This is the instruction that resulted in the error. I usually disassemble them by hand and see what it is, then I try to find out exactly where in the function this assembly sequence comes up (probably by compiling the .c file to a .s file and looking around). Alternatively you can just run `gdb` on `/usr/src/linux/tools/system`, but `gdb` has some problems.

- `task[0] (swapper) killed: unable to recover`
`kernel panic: Trying to free up swapper memory space`
`In swapper task - not syncing`

This is just to tell you not to expect anything from the kernel any more: linux cannot go on without the swapper task.

9.2.23 How can I decode the numbers in the message?

Also, note that sometimes the `Unable to handle...` message is followed by the `Oops: 0002 block`, but the last message is sometimes something else (i.e. not always a `task[0] (swapper)...` message, but something from the TCP driver in this case).

It depends on where it happened. If it's an interrupt, it can essentially happen in any task (all interrupts are handled in the task that happens to be currently active), while *normal* code happens in the task that calls it.

10 Emacs for Linux

This section is maintained by Rick Sladkey (jrs@world.std.com)

Last Update March 1993.

10.1 General Information

This is a short list of Frequently Asked Questions about GNU Emacs under Linux. It does not address general questions about Emacs which are not Linux specific. For general help about Emacs,

1. learn to use and read the online documentation
2. read the real Emacs FAQ found in `emacs/etc/FAQ` and
3. read the newsgroup `gnu.emacs.help`

10.2 GNU Emacs for Linux

10.2.1 Does Linux Emacs use the shared libraries?

Yes. It wasn't easy. If you have any problems with them let me know. Also contact me if you would like static binaries.

It works fine with the DLL libraries and should not require a new binary when the C or X libraries are updated.

10.2.2 What version of the compiler was used? Which shared libraries?

This describes version of GNU Emacs 18.59 for Linux 0.97 and above compiled with GCC 2.2.2d7 using the `libc.so.4.1` and `libX11.so.2.1` shared libraries.

10.2.3 Which files do I need?

<code>emacs-18.59a.tar.Z</code>	this file, sample <code>default.el</code> and diffs for this version, the eight-bit patch, <code>iso-latin-1.el</code> and <code>eight-bit.el</code>
<code>emacs-etc-18.59a.tar.Z</code>	emacs support programs and misc info
<code>emacs-bin-18.59a.tar.Z</code>	shared emacs binary and its doc file
<code>x11emacs-bin-18.59a.tar.Z</code>	shared x11emacs binary and its doc file

10.2.4 How do I install them?

According to Linux convention, cd to `/usr` and untar the packages.

10.2.5 What if I want to compile Emacs myself?

It is fairly easy (and highly recommended) to compile Emacs yourself if you have the the disk space. In this case you only need the standard Emacs distribution (`emacs-18.59.tar.Z` from any GNU archive) and the Linux diffs and support files (`emacs-18.59b.tar.Z` from a Linux archive). The diffs are quite small and mostly amount to a configuration file.

10.2.6 What about Epoch or Lucid Emacs? Are these available for Linux?

Yes. Thomas Dunbar has been maintaining Epoch for Linux and Chipsy Sperber has compiled Lucid Emacs. Both of these work well under Linux. Look in a Linux archive index for where to find them.

10.2.7 Does Linux Emacs support eight-bit input/output?

Yes. It is new with this version. It uses the so-called *ctl-arrow* patch. See the file `README.8bit` for more information.

10.2.8 How much disk space is required?

Anywhere from 1 to 15 Meg. Emacs works reasonably well with no support files at all. With a judicious selection from `lisp/*.elc` and `etc/*`, quite a lot can be done using only 2 Meg. If you want all of `lisp/*.elc`, `info/*`, and `etc/*` this will require 4 to 6 Meg. If you unpack the whole source you need 8 Meg. If you collect info files like rare coins and install a lot of big lisp packages then Emacs may need its own partition. :-)

10.2.9 Why can't Emacs find its support files anymore?

This is because older versions of Emacs were compiled with `/usr/local/emacs` based paths. The current version is compiled with `/usr/emacs` paths. If you have a previous installation, just `mv /usr/local/emacs /usr` and you're done. If you can't bear to part with the `/usr/local` pathnames because of inertia then do `ln -s /usr/local/emacs /usr` and you can have them both.

10.2.10 How do I get Emacs to recognize my cursor keys?

Simple. Don't use them. :-) Seriously, there are as many ways to do this as there are elisp hackers but the preferred way is to follow the pattern set by the other terminal definition files in `emacs/lisp/term/*.el`. For just arrow keys you can just copy `vt220.el` to `console.el` and that's it. For function keys and the others see the sample `default.el` included with this file.

10.2.11 What packages are particularly useful under Linux?

Because info format is the documentation standard of the GNU project and just about everything except the kernel comes from FSF, you will find that Dave Gillespie's enhanced info package is very useful. It allows multiple info directories, space bar paging, and supports compressed info files. Please learn to use info. Imagine Unix life without man.

Others that I highly recommend are Sebastian Kremer's enhanced dired directory editor, Dave Gillespie's calc calculator, Masanobu UMEDA's gnus for usenet news, and Kyle Jone's vm for mail.

All can be found in the OSU Emacs archive. See the real FAQ for more details.

10.2.12 Does Linux Emacs support the X Window System?

Yes. However, there are two binaries. One without X support (about 485k) and one with X support (about 515k).

10.2.13 Do I need both Emacs if I don't always use X?

No. The X11 version works equally well inside or outside of X. If you get the message `Check your DISPLAY variable` it means that you have defined `DISPLAY` in your `~/.profile` (or whatever). You can fix this by starting Emacs with `emacs -nw` or by removing the `DISPLAY` variable from your `~/.profile` and putting it in your `~/.xinitrc`.

10.2.14 Why doesn't Emacs use the settings in my `.Xdefaults/ .Xresources` file?

You are probably using the word `emacs` and your X version of emacs is called `x11emacs`. Either use the word `Emacs` in your resource file or rename `x11emacs` to `emacs`. See above question on why this is reasonable.

10.2.15 I read about some menu that is supposed to pop up when I press some mouse button. Does this work with Linux Emacs?

Yes. This requires XMenu support to be compiled in. Former versions did not support it because it did not work correctly.

10.2.16 Sometimes Emacs crashes with a `SIGALRM` message. What's wrong?

You need to upgrade to a newer version of `bash`. Some older versions of `bash` were compiled with buggy libraries that caused `bash` to send a `SIGALRM` to its parent in some circumstances. Recent versions do not exhibit this problem.

11 Features

11.1 Virtual Consoles

11.1.1 I've read that Linux has virtual consoles, what must I do to get them?

Yes there are, you can access them with the left `[alt]`-key together with `[Fn]`-key. With the Linux Images distribution, 4 consoles are available, `getty` runs on them.

Notice that they are NOT accessible when running X (contrary to some commercial unices).

11.1.2 When Linux boots, I get the following message "8 virtual consoles"; how can I access to the 5-8 vc's?

If you want the `getty` to run on the 5-8, you should add the corresponding entries in `/etc/inittab`. You can also just run `sh` on them by using the `doshell` soft. And then in either case, the `ALT-F[5-8]` will access the corresponding `vc`.

11.2 Miscellaneous

11.2.1 What kind of shell is /bin/sh?

It's the Bourne Again Shell, `bash-1.12.3` and compilation was straightforward, just `make` that's all or nearly.

BTW There does exist different shells for Linux, these are: `bash`, `rc`, `zsh`, `tcsh` and `pksh` (a korn shell).

11.2.2 Does there exist a man page for ****?

Download `man.tar.Z` from your favorite Linux ftp site, there is most of the fileutils man page — either `****` or `g****`, example there is nothing on `ld`, but there is for `gld` :) —, check the `whatis` database provided. The files in the `cat1` dir are pre-formatted man pages that the man program can use. Quite recently the man pages for section 2 have been written (thanks Drew) and can be found, at least at `tsx-11` in `/pub/linux/docs/man/man2.tar.Z`

Also manpages are in the SLS on the b? disks.

BTW there is `nroff` and `groff` for Linux. `Cawf 2.0` works just fine for simple man pages, and a partial ms support too.

Moreover Michael Johnson is the coordinator for man pages under Linux, he is looking for volunteers, so contact the DOC Channel.

11.2.3 Is there a simple man package (groff is too big)

`Cawf -man` also `fm -m` (Al Clark) `tsx-11` in `/pub/linux/binaries/usr.bin/fm.tar.Z`

11.2.4 I have a 2MB box, but free reports only 1MB why?

`free` doesn't concern with the memory the kernel has allocated for itself. In other words what you see is the `user` memory available. The kernel has taken the low 1Meg for it's use (~250Ko for it and the rest for buffer cache and kernel data structures); for big memory machine it could be even 2MB.

11.2.5 Does there exist shared libs?

H.J. Lu, (`hlu@eecs.wsu.edu`) writes on 09/01/92:

The shared library under Linux started at 0.12. Peter MacDonald collaborating with Linus made the first generation of shared library, which is the base of the current classic shared library.

The kernel support of shared library under Linux is system call

```
extern int uselib (const char *__filename);
```

which loads an executable image with fixed entry point into memory, just like the ordinary executables.

In `crt0.s`, a function which can find out if and which shared images are needed and loads them is invoked before `main ()` is called if necessary. David Engel and I developed a way to tell the loader which shared images have to be loaded, utilizing the similar technique used in global constructor in g++ 2.x with the help from GNU binary utilities.

In the classic Linux shared library, we build a big executable image for several libraries and make sure no external variables outside of the participating libraries are referenced. Then we can get the absolute addresses of all the global variables defined in the libraries used to build that executable image. After that, we make a

stub library for each participating library which just has the absolute addresses of all the global variable in it.

For each shared image, there must be one and only one file, usually called, `__shared.o`, which defines a global variable containing version, name and entry point of the shared image, and a dummy global data. Among those libraries used to build the shared image, there must be one library which will always be referenced whenever any other library is referenced. We put `__shared.o` into the stub library for that library and add a declaration for the dummy global data defined in `__shared.o` which will make sure `__shared.o` will always be linked in when any participating libraries are linked.

In gcc 2.2.2d, jump table, developed by David Engel, was introduced in the shared library. At the beginning of each shared image, there is a table in which every library function has a fixed entry address and the instruction at that address is a jump which will lead to the real library function. So we can change the library function without changing the corresponding entry address of the jump table. For the global data we put them at the beginning of data section of the shared image. We have to separate them from text code and link them in fixed order. It is very hard to maintain the same addresses for the global data when library is changed. After the global data are set up properly and some spaces are left for possible future changes (that is a very tough procedure.), it isn't too difficult to maintain.

In the current implementation, only `libc.a`, `libcurses.a`, `libdbm.a`, `libtermcap.a` and `libm.a` are built with jump table. The global data in X11 libraries are too complicated to make jump table such that their addresses won't change when there is a change in X11 libraries. It's not apparent yet that the benefits gained from a jump table version of the X libraries would offset the effort required to set it up and maintain it unless we get some cooperation from X Consortium, which is very unlikely. But they are linked with jump table version of `libc.a` and `libm.a`. That means they don't have to be relinked when there is a modification in `libc.a` or `libm.a`.

11.2.6 Why do I need dynamic linking?

The best way to illustrate this is with a simple example of what happens without dynamic linking.

```
#include <stdio.h>

int errno;
FILE * foo;

main(){
foo = fopen("/usr/bin/foo","w");
printf("foo: %x errno: %d\n",foo, errno);
}
```

Assume this program is run by a non-root user. Normally you would expect that the open will fail because the user does not have permission to write to `/usr/bin` — this means that `fopen()` would return NULL, and `errno` would contain 13, which means EACCES-Permission denied. If you link this to a normal jump table library, it prints

```
foo: 0  errno: 0
```

Now consider the following nearly identical program:

```
#include <stdio.h>

extern int errno;
FILE * foo;

main(){
foo = fopen("/usr/bin/foo","w");
printf("foo: %x errno: %d\n",foo, errno);
}
```

The results for this program are:

```
foo: 0  errno: 13
```

Why does this happen???. First of all you must realize that in the first program `errno` is not declared external, and thus storage will be allocated in the data segment for the variable. Unfortunately there is no way to tell libc this, and libc has it's own storage for `errno`. Thus libc writes the result in one location and you are looking in the other. In the second program `errno` is declared extern, and thus no storage will be allocated. The linker will see the need for a definition of `errno`, and it will find one in libc. Thus in the second example the user program is reading the same memory location that the library stored the answer in.

This is of course a simple example. There are more complicated examples involving programs that want to redefine a library function and these can lead to quite bizarre behavior which is difficult to debug.

With dynamic linking, both of the above programs yield the same (correct) result.

11.2.7 How does dynamic linking work under Linux?

Basically, to implement dynamic linking we need to route all function calls to global functions through the jump table. We also need to make sure that all global data is accessed indirectly through a pointer. The actual libraries do not really look very much different from the actual jump table libraries that we are currently using, but there are additional pointers to the global data in the sharable image that are used internally by the library to access all global data. There are also additional symbols in the stub library that are defined as the address of the various pointers, and are used by the linker to help determine whether there is a pointer or a jmp instruction that needs to be fixed up.

The linker has one new important task with dynamic linking. It watches for duplicate definitions of symbols, and if the second definition is in a sharable library, it makes a note of it as a conflict that needs to be resolved at run time. The linker puts a list of all of these conflicts into a table and makes it available to `crt0` through the variable `__SHARABLE_CONFLICTS_`.

At run time, `crt0` (through the function `__dynamic_resolve`) will go through and fix up all of the pointers that need to be fixed. Let us say for example that you define your own `malloc` function in a program — in this case at run time the startup code will actually modify the jmp instruction in the jump table so that it points to the `malloc` function in your program, not the one in the library. The library will never refer to the `malloc` function directly, but it will always be referenced through the jump table — in this way we ensure that there will only be one function named `malloc` that is used by a given program.

If you were to run the first example program in the "Why do I need dynamic linking" answer, then the pointer in the library that corresponds to the variable `errno` will be modified so that it points to the `errno` variable that is in your program.

The library itself will never use `errno` directly — it will always use it via the special pointer variable. In this way we ensure that all references to the variable `errno` both in your program and in the sharable library will reference the same memory address.

11.3 Swapping

11.3.1 How do I make swapping work?

Quite simply, you need the `swapon` and the `mkswap` binaries. Then you can choose between a swap partition or a swap file. The `mkswap` is used to write the *swap signature*, while the `swapon` binary is to activate the swapping.

First of all you need a partition :), I assume it's the second of your first disk namely `/dev/hda2`, and it's 10MB big

1. swap partition

you have to indicate it's a swap area, this is done via `mkswap` (instead of `mkfs`) which needs the name of the partition and the size in blocks (a block is 1Ko big); the optional `-c` flag is for bad block checking. So for our example you should perform:

```
mkswap [-c] /dev/hda2 10000
```

Then you need to indicate that you want Linux to use the swap area, this is done via `swapon`. In general it is set in the `/etc/rc` file, just put the following entry:

```
/bin/swapon /dev/hda2
```

It can also be achieved via the `/etc/fstab` file

2. swap file:

The process is quite close; you need a partition, and a swap file. Assume that I prefer a swap area of 4MB (I want to keep some place in `/dev/hda2`). I need first to `dd` the file.

```
dd if=/dev/hda2 of=/swap_file bs=1024 count=4096
```

`bs` stands for block size, and `count` is the number of blocks then I have to put the *swap signature* on that file:

```
mkswap /swap_file 4096
```

At this point, you should `sync`, just to be sure the signature is effective. And finally add an entry in the `rc` file:

```
/bin/swapon /swap_file
```

11.3.2 When I boot I get one of the following messages: Unable to find swap signature or Bad swap-space bitmap

You probably forgot to make your swap-device, use the `mkswap` command.

11.3.3 How do I know if it is swapping?

You will notice it :)) First of all, Linux tells you at boot time, **Adding swap: XXX pages of swap space**, and if you start running out of memory, you will notice that the disk will work overtime, and things slow down. Generally a 2Meg RAM will make the system swap constantly while running gcc, 4 Meg will swap occasionally when optimizing big files (and having other things active, such as make). Also, the command **free** (from the **ps** package) reports total enabled swap space and current swap use.

11.3.4 How is it possible to remove a swap file?

Simply perform a **rm** on that file, and remove the **swapon** of your **/etc/rc** file.

11.3.5 How is it possible to remove a swap device?

mkfs the device, and remove the **swapon** of your **/etc/rc** file.

11.3.6 How much swap space do I need?

Linux does not perform real swapping, it's rather paging (see below for a more complete explanation). The swap area is *added* to the memory and can be viewed as virtual memory, so choose the size you need, example:

8MB RAM + 6MB swap \Rightarrow 14MB virtual memory

11.3.7 Could someone explain the swap process on Linux? Is it swapping or paging?

Linux uses swap as *additional* memory, one page of the swap-space is used for the good-page bitmap and the swap-space signature.

In fact Linux does only paging, no swapping in the meaning *write out one whole process to disk*.

The reason it's called swapping is that Linux used paging for memory management on a low level since the very beginning, but didn't page to disk at all until 0.12.

11.3.8 Is demand paging different from paging and how?

Demand-paging is really *demand loading of executables* and is totally independent of the page-swapping algorithms, although they have similarities. When Linux starts up a process, no actual code space is loaded: I let the page exceptions load in the executable as needed. Thus Linux demand-loads the code and initialized data it needs.

Demand-loading has very good points:

1. it simplifies the exec system call
2. it means page sharing between that have executed the same file is easy to implement
3. it cuts down on the amount of memory required.

When Linux runs out of real memory, it starts to look for pages it can swap out, but if it notices that the page is clean, it just forgets about it, and demand-loads it when it's needed again. That means that swap-file isn't needed as much, especially

when running big binaries such as gcc, where the code-pages can be demand-loaded as you wish.

Point 3 means that even without any swap space, you can usually run slightly larger programs that your memory setup would actually permit. I've noticed this while running X and doing a kernel compilation + something else when I've forgotten to turn on swapping: free reports 0 page available but things still work, although performance is slightly down...

11.3.9 Is there any way to tell how much swap space you are using or have left?

The free program provided with the ps package handles this.

11.4 Patching

11.4.1 How does one go about applying a patch to Linux?

In the UNIX world most of distribution are in source form. This includes the operating system. To apply a patch, you apply it with the patch program to the affected sources. The patch program takes as input the differences between the old and the new version. After patching you need to recompile the sources.

Assume I want to apply a patch enclosed in the file XXX. First of all I will look at the top of XXX, where the file affected is identified. This may have a leading path attached to it. Either cd out to the root of the patch, ie if I see linux/kernel/blk_drv/blk.h I would cd into /usr/src (assuming it's the place where I can find linux/kernel...) and then patch as follows

```
patch -p0 < whatever_place/XXX
```

or, you can specify a number of path components to strip from the path. If I am in the blk_drv directory patching would be

```
patch -p3 < whatever_place/XXX
```

11.4.2 There are a lot of patches available (ps patch, NFS patches, CD-ROM patches ...) can I be fairly confident the subsequent patches will work?

This is not true yet for the current version; but it will be so I kept it :)

No you can't, patching is a real beta tester art :)). People are not working on the same patched release, so you have to check if the patches you already applied works on the same kernel part, if not, *great*, just apply them. If yes, check if there is an order, patch creator knows that, and (should) try to warn patch user (in other words: beta tester) otherwise you should edit the patch files (and possibly make a brief note to others on this list/newsgroup or even a *cdiff*) before applying them, another solution is to keep cool and wait for the next version of Linux where, in general, the modifications have been done but this behavior is *not* Linux helpful.

11.4.3 I got the patches on some ftp sites, and applied them to the kernel and tried to compile. It didn't !!. Are the patches buggy?

Before remake, just do a make clean in the directories involved by the patches. This will force a rebuild of the .o and .a files.

If you have a RCS running on your source tree, did you checked a patched version of the files changed before *any* CO either by you or make

Finally, make sure the patches succeeded. Normally, failed patches on a file **FILE** will leave a **FILE#** file. Moreover you will get a **chunk failed** message. It is possible to capture the output while patching, with the following:

```
patch -p0 < patchfile | 2>&1 patch.result | more
```

11.5 Software

11.5.1 What are the editors available in Linux?

Right now there are uemacs-3.11, elvis-1.4 (1.5). GNU Emacs is there, read the section 10. Also the port of mg (micro gnu) has been done and can be found at least at athos.rutgers.edu (128.6.4.4) in `pub/linux`, mg is the binary and `mg.tar.Z` is the sources file. You can also find a PD ed, and elvis has an ex mode. Finnally there are joe, vile-3.11, elle (Elle Looks Like Emacs), Xedit and aXe.

11.5.2 Does there exist a printer package for Linux?

Yes. The `plp` package is currently available under the directory `/pub/linux/BETA/plp` on `tsx-11` and its mirrors. You may also print things manually like so:

```
cat filename > /dev/lp1
```

(Note that though `/dev/lp0` exists, most people find that their printer is on `/dev/lp1`. Use whatever the kernel says that it detects in the boot-up messages.)

11.5.3 Are uucp, mail, and/or USENET news available for Linux?

Yes !!! At this time, the following (at least) have been ported to Linux:

Package name	Version
taylor uucp	1.03
elm	2.3-pl11
smail	3.1.28
C-news	12/22/91
tin	1.1pl4
trn	2.2
nn	6.4.18

[...coming soon - Expect an announcement in the near future concerning the packages above being made available in binary form in the SLS distribution and in source form from sunsite.unc.edu and other Linux archive sites...]

If you are interested in the last uucp features, join the UUCP channel (see section 4.1.4 for more details about the multi channels list).

11.5.4 What is VFS?

Linux 0.96 already has Virtual FileSystem, which means that it acts as a filesystem switch. It makes it easy for someone to design another filesystem format and include it in the Linux kernel along with the standard minix filesystem format. So it *enables* someone to design a robust filesystem which would have some nice properties (no 14 chars file name limitation, nor 64Meg limit), and could be included in the kernel in such a way that both the Minix and the new one could be mounted at the same time. This solves the uncompatibility problem; since the root disk could still use the Minix filesystem, while the hardisk could be using the new one.

11.5.5 What's about T_EX?

The primary site for Linux T_EX is 129.78.66.1, this is P. Williams' site in Australia. The stuff at `tsx-11` was posted by T. Dunbar who does support/maintain the `dvilj` stuff.

11.5.6 What's about LILO?

Werner Almesberger, the creator writes about LILO, a Generic Boot Loader for Linux (*Linux LOader*)

This is an ALPHA test release of a new boot loader. Be sure to have some means to boot your system from a different media if you install LILO on your hard disk.

- Features
 - does not depend on the file system. (Tested with Minix, EXT FS and MS-DOS FS.)
 - can be used to boot from floppies and from hard disks.
 - can replace the master boot record.
 - can boot non-Linux systems (MS-DOS, DR DOS, OS/2, ...) and unstripped kernels. item- supports up to 16 different boot images that can be selected at boot time. Root and swap disk/partition can be set independently for each image.
 - boot sector, file map and boot images can be all on different disks or partitions.
- Restrictions and known problems
 - SCSI disks are not fully supported yet. (Still waiting for some kernel changes.)
 - booting other operating systems doesn't seem to work everywhere. If everything but booting a non-Linux OS from LILO works on your system, you should boot LILO by `BOOTACTV` and select the alternate OS with the latter as a temporary work-around.
 - booting non-Linux systems from the second hard disk ("D:") is not yet supported.

Please send all bug reports to `almesber@nessie.cs.id.ethz.ch`.

11.5.7 What's about MGR?

There is a MGR channel available, contact the request adress with help in the body: `linux-activists-request@niksula.hut.fi` The stuff can be found at `tsx-11` in `pub/linux/packages/MGR`. In brief:

MGR provides:

- multiple overlapping windows
- multiple fonts
- text and graphics in each windows
- a simple popup menu package
- a client/server model *a la* X

- independence from any peculiar networking technology

MGR consist of a server process and some clients. Each client has his own window, and can create subwindows. Clients communicate with the server via a bidirectionnal channel. A C library is provided.

When a new window starts, it is as a terminal emulator running the shell; for more information you can grab the `mgr-man.out` from `bellcore.com`

The information hereafter appeared on the MGR channel on 19th Jan.

Unpack the `src*` stuff under in `/src/lbin/mgr` and the `usr*` stuff in `/usr/mgr`. I splitted things that way for making it easier to ftp them and carry them home. If you don't like these paths, then use others, but remember to edit the Configfile after generating it. **YOU NEED TO INSTALL ALL FILES FOR COMPILING.**

This is a beta release for programmers, there are no binaries. It makes my work available for people who either also work on MGR and like to get new sources and for advenurous beta testers. **WARNING:** I assume using Linux 0.99.2 and GCC 2.3.3.

There is some new m4 code to generate menus. It is more powerful than `menu(1)`, but a little harder to use.

MGR is currently being ported to Coherent 4.0, but this version doesn't include everything yet. There is already a beta version running on Coherent, although not too stable yet.

For questions, write to the mailing list, channel MGR. Please be sure to talk about the latest version of MGR plus having read the README files in the directories containing the problem.

You got this stuff from `ftp.thp.uni-koeln.de:pub/linux/mgr` or from `tsx-11.mit.edu:/pub/linux/packages/MGR`, which is a mirror of the Cologne directory.

11.5.8 I have successfully compiled MGR, but when I try to run the program I get can't find mouse or already in use, any clue?

Try the following `mgr -m /dev/ttys1` if the mouse is on the serial 1. Another possibility is to link `/dev/mouse` with `/dev/ttys1` (assuming your mouse is on serial 1). Or if it's a bus mouse, `mknod /dev/mouse c 10 0` once.

11.5.9 Any tips for MGR?

Well, I have tried it on my 386SX EGA/VGA; the screen is Ok but the Logitech mouse I have is not well recognized.

BTW check the major/minor number for pty's; they should be character device with 4 as major and 128 and bigger as minor:

```
ptyp0 c 4 128
ptyp1 c 4 129
...
ttyp0 c 4 192
ttyp1 c 4 193
```

11.5.10 What's about X11?

See the section 14 devoted to X11 in this FAQ.

11.6 Hardware

11.6.1 What tape drives work with Linux?

- There is a working QIC-02 device driver for Linux, supporting (at least) Everex/Wangtek cards.
- There are additional patches for the QIC-02 to support Archive SC402/499R. You can find them in `/pub/linux/alpha/qic-02` at `tsx-11`. There are some bugs in the driver but you can backup.
- Newer drivers are all SCSI, so check the SCSI section (13) in this FAQ.

11.6.2 Is there only the *** keyboard?

There are Dannish, Finnish, French, German, Uk, US and DVORAK keyboards. Set it in the main `Makefile` of the kernel sources, then (re)compile the kernel again. Make sure the files in `kernel/chr_drv` directory are recompiled.

11.6.3 (special FINNISH/US) I booted up with the new image and everything work except that some keyboard keys produce wrong characters. Does anyone know what is happening?

Since 0.95a images are US product (and so are US-keyboard oriented), BUT Linux sources are FINNISH product, and so the default keyboard is set to be FINNISH. The solution is in the previous section.

11.6.4 Does Linux work for SCSI drives?

Yes since v0.96. At `tsx-11.mit.edu` in `/pub/linux/ALPHA/scsi`, you will find the last SCSI alpha/beta version and also a special SCSI FAQ, read it, it contains the latter information than the one provided in the SCSI section (13). You should, also, contact the `linux-scsi` list or directly `drew@cs.colorado.edu`.

11.6.5 Linux is supposed to work with ESDI drive. However I have trouble with my Magtron MT-4115E (Joincom controler), any clue?

Some harddisk don't like Linux (even though they should). Maybe not a bug but a deficiency.

I had to remove the `printk unexpected hd interrupt` statement in `hd.c` because I was getting so many of those messages. Be warned that if there is any read error the system just hangs, even the `ctrl-alt-del` won't work. You should be able to use your ESDI drives if you could live with those nuisances.

11.6.6 What's about Bus Mice?

Since the Linux v0.96c-pl2 the kernel does support LOGITECH and BUS MICE If you are unsure that you have a bus mouse or not, check to see if your mouse card has a selection for a sample rate switchable between 30Hz and 60Hz (or possibly 25/50Hz), if it does not, then it is NOT a true bus mouse (InPort mice for example will not work with this driver).

To create a bus mouse device:

```
mknod /dev/mouse c 10 0
```

12 GCC Information

12.1 HLU Information

This section of the FAQ is maintained by H.J. Lu (hlu@eecs.wsu.edu). Please mail any corrections or changes to him.

Last update March 1993

12.1.1 I don't know how to install gcc stuff, is there special places?

This is the `README.gccdisk`

- Introduction

This is GNU C/C++ 2.3.3. It is linked with jump table 4.2. You should use it with the Linux C library disks since the C/C++ compiler has no shared images, which are required by all the binaries in it. It consists of 2 minix disks, which will fit on either 5.25" or 3.5" floppies.

It is on `tsx-11.mit.edu` under `pub/linux/GCC/gccdisk`

- Distribution File Format

There are two files, each of which goes on one disk

1. Disk 1 (`gcc233a.Z`)

This contains `gcc`, `cpp`, `cc1` and `crt0.o/ gcrt0.o`. There are also some header files in `/usr/lib/gcc-lib/i386-linux/2.3.3/include`.

2. Disk 2 (`gcc233b.Z`)

This disk contains `cc1plus`.

- Installing on the Hard Drive

1. uncompress `base[1|2].Z`.
2. rawrite or dd each file to a formatted floppy disk.

Now you have made a copy of gcc 2.3.3 on two floppies. To copy each floppy to your hd, you should mount the floppy and copy its contents to your hard drive. You can do this by:

```
mount /dev/fd[0|1] /mnt
cd /mnt
for d in bin dev etc usr
do
if [ -d $d ]; then
cp -av $d /
fi
done
```

Note: This may overwrite some files on your hard disk.

12.1.2 What compiler should I use for Linux?

You should only use the same version on `tsx-11.mit.edu` or `fgb1.fgb.mw.tu-muenchen.de` under `/pub/linux/GCC`. If you want to use the testing release, first join the GCC channel on the Linux mailing list (look section 4.1.4), and then send a note to hlu@eecs.wsu.edu. Don't use gcc older than the one on `tsx-11.mit.edu` or `fgb1.fgb.mw.tu-muenchen.de`.

12.1.3 Where is the latest official gcc 2.xx for Linux?

It's on `tsx-11.mit.edu` under `/pub/Linux/GCC` and `fgb1.fgb.mw.tu-muenchen.de` under `/pub/linux/GCC`. You may find it on the other sites. Since gcc 2.3.3, you can compile it yourself out of box from any gnu ftp sites. Just unpack the source code and do

```
configure [i386-linux|i486-linux]
```

You should follow the instructions in `INSTALL`.

12.1.4 Where is the latest official Linux C library?

It's on `tsx-11.mit.edu` under `/pub/Linux/GCC` and `fgb1.fgb.mw.tu-muenchen.de` under `/pub/linux/GCC`. You may find it on the other sites.

12.1.5 What are the contents of them?

Please read the current release note and `ChangeLog` for details.

12.1.6 How do I install it?

Read `README` and release notes.

12.1.7 What are the main differences with the old release?

Read `README` and release notes.

12.1.8 Can I use the old version of gcc?

Please get rid of gcc older than gcc 2.2.2. Starting from gcc 2.2.2, you can do

```
gcc -V xxxx
```

where `xxxx` is the version number. Please read `release.xxxx` for detail. There is one catch in gcc 2.2.2d, `setjmp/longjmp` is changed, so the old header files is not compatible with gcc 2.2.2. Before you install gcc 2.2.2d, please do

```
cp /usr/include/setjmp.h /usr/lib/gcc-lib/i386-linux/2.2.2/include
```

where `/usr/include/setjmp.h` is come with gcc 2.2.2.

12.1.9 Can I delete the old shared image in /lib?

Since the Linux C library version 4.3, we have introduced the dynamically linked library. There is no need to keep old shared images in `/lib` if none of your binaries are linked with the classic shared library which has been obsolete. You can just keep one version of the shared image in `/lib` which has the highest major and minor release numbers for each shared library. The name of the shared image is `xxxxxxxxxxxxxx.so.major.minor`.

12.1.10 Is stdio ANSI compatible?

Yes, please test it.

12.1.11 Is g++ in 2.xx?

Yes.

12.1.12 Where can I get the gcc manual?

You can get man pages and manual come with the gcc source code on any GNU ftp sites. You should find `gcc-man.tar.z` on `tsx-11` which has man pages for gcc.

Read the manual page, `gcc.ps` or `gcc.man` in `/usr/install/gcc2`.

12.1.13 What options can I use for gcc?

Read manual page, `gcc.ps` or `gcc.man`. Also `-static` tells gcc to use the static libraries, `-nojump` forces gcc to use the classic shared libraries. The default is the jump table version of shared libraries. The shared libraries for X are linked with the jump table version of shared C library.

12.1.14 How can I debug the C code?

Read the gcc/gdb manuals. The simple way to do is

```
gcc -g foo foo.c
gdb foo
```

12.1.15 Where is the source code of the new libc.a?

The same place you find this file. It is called `lib-src-yy.xx.TZ`.

12.1.16 Why does g++ complain, even die?

You need `expr`, which is in GNU shell utilities 1.6, `echo (?)` and `sed`.

12.1.17 How do I generate code for i486?

Add `-m486` to `CFLAGS`.

12.1.18 I heard malloc(0) wouldn't work with Linux, what should I do?

include `<stdlib.h>` and don't define `NO_FIX_MALLOC`.

12.1.19 Why does gcc say "xxxxx..h not found"?

See section 12.1.5 (What are the contents of them?).

12.1.20 I really followed every step in the documentation, but when I do make, why does it say don't how to make xxxxxx?

The dependency in `Makefile` is dated, you need to make a new one. Please get some guide on make and read `Makefile`. For the kernel sources, please do

```
cd /src/linux
make dep
```

12.1.21 How do I compile programs under Linux?

The Linux C library is trying to be ANSI/POSIX compliant. It is also very compatible with SYSV and BSD. The C library is loaded with SYSV and BSD functions. There are three exceptions:

1. signal in Linux is POSIX.
2. tty in Linux is POSIX.

3. time functions are POSIX, plus a few BSD and SYSV extensions.
4. `setjmp/longjmp` functions are POSIX. But you can use `-D__FAVOR_BSD` to make it BSD or use `sigsetjmp/siglongjmp`.

When you compile a program under Linux, your best bet is include all the appropriate header files and use `-Wall`. All the usable functions and global variables are declared in the corresponding header files. **YOU SHOULD NOT DEFINE ANY functions or global variables OF THE LINUX C LIBRARY IN YOUR CODE IF YOU WANT TO USE THE SHARED LIBRARIES.**

After saying all those, you now should know you can compile a program with `-D_POSIX_SOURCE` or `-D_GNU_SOURCE` (read `<features.h>` for details). With a few modifications you can even use `-DSYSV`, `-DUSG` or `-DBSD`. Some codes need to define `-DSTDC_HEADERS` for ANSI C compiler like gcc here.

To use `malloc()` and `calloc()` safely under Linux, please include `<stdlib.h>` and don't define `NO_FIX_MALLOC`.

BTW, gcc `-traditional` should work with gcc 2.2.2d or above.

Please also read `ChangeLog` for the latest enhancement.

Please read the header files for details. Maybe you should get a book on POSIX. Any suggestion of the book list? Refer to the ultimate list in section 1

I like *POSIX Programmer's Guide*, by Donald Lewine. Its essentially a list of POSIX functions' man pages, with a very brief guide in the beginning of a few things. It's published by O'Reilly & Associates, Inc. (from Steve Robbins (steve@nyongwa.cam.org))

12.1.22 How can I get bsd style signal?

Use `-D__USE_BSD_SIGNAL`.

12.1.23 Why does a program that should only poll for input become a CPU hog?

The `select()` system call. The timeout parameter was classically used read-only by the system. Some manual pages already notes three years ago:

`select()` should probably return the time remaining from the original timeout, if any, by modifying the time value in place. This may be implemented in future versions of the system. Thus, it is unwise to assume that the timeout pointer will be unmodified by the `select()` call.

If you do not take this advice seriously you get a zero timeout written back to your timeout structure, which means that future calls to `select()` using the same timeout structure will immediately return.

Fix: Put the timeout value into that structure every time you call `select()`. Change code like

```
struct timeval timeout;
timeout.tv_sec = 1; timeout.tv_usec = 0;
while (some_condition)
{
    select(n, readfds, writefds, exceptfds, &timeout);
}
```

to

```

struct timeval timeout;
while (some_condition)
{
    timeout.tv_sec = 1; timeout.tv_usec = 0;
    select(n,readfds,writefds,exceptfds,&timeout);
}

```

12.1.24 When a program is stopped using Ctrl-Z and then restarted, or in other situations that generate signals: Ctrl-C interruption, termination of a child process etc. why does it complain about "interrupted system call" or "write: unknown error" or things like that?

The system call the program was executing has been interrupted to process the signal, and then it returned -1 and set `errno = EINTR`. The program then was likely to draw bad conclusions from that.

Explanation: Your program has signal handlers installed, using `signal()` or `sigaction()`. When the signal occurred, your signal handler was invoked. In other UNIX systems, this usually happens asynchronously or in a few slow system calls:

When a signal is caught during the execution of system calls such as `read(2)`, `write(2)`, `open(2)` or `ioctl(2)` on a slow device (such as a terminal, but not a file), during a `pause(2)` system call or a `wait(2)` system call that does not return immediately because a previously stopped or zombie process already exists, the signal-catching function is executed and the interrupted system call then returns a -1 to the calling process with `errno` set to `EINTR`.

Linux (following POSIX) checks for signals and may execute signal handlers

- asynchronously (at a timer tick),
- on return from *any* system call,
- during the execution of the following system calls:
 - `select()`,
 - `pause()`,
 - `connect()`, `accept()`,
 - `read()` on terminals or sockets or pipes or `/proc` files,
 - `write()` on terminals or sockets or pipes or line printer,
 - `open()` on FIFOs or PTYs or serial lines,
 - `ioctl()` on terminals,
 - `fcntl()` with command `F_SETLKW`,
 - `wait4()`,
 - `syslog()`,
 - any TCP or NFS operations.

[For other operating systems you may have to include the system calls `creat()`, `close()`, `getmsg()`, `putmsg()`, `msgrcv()`, `msgsnd()`, `recv()`, `send()`, "wait()", `waitpid()`, `wait3()`, `tcdrain()`, `sigpause()`, `semop()` to this list.]

In the last two cases and assuming the program's signal handler returns, the system call returns -1 and sets `errno` to `EINTR`.

If the `SA_INTERRUPT` flag is not set for the corresponding signal, however, in most cases the system call is automatically restarted (continued) after execution of the signal handler, and your program won't see any `EINTR`.

You may ask why this is not the default behavior when the default Linux `signal()` is used to install the signal handler. This is because POSIX adopted this. As for which one is better, it is a matter of opinion.

Note that in some versions of BSD UNIX the default behavior is to restart system calls. To get system calls interrupted you have to use the `SA_INTERRUPT` flag.

Fix: Either add `-D__USE_BSD_SIGNAL` to your `CFLAGS`. Or for every signal handler that you install with `signal()`, use `sigaction()` instead, without setting `SA_INTERRUPT`.

Note that while this applies to most system calls, you must still check for `EINTR` on `read()`, `write()`, `ioctl()`, `select()`, `pause()`, `connect()`. You may do it like this:

```
int result;
while (len > 0)
{
    result = read(fd,buffer,len);
    if (result < 0) break;
    buffer += result; len -= result;
}

→

int result;
while (len > 0)
{
    result = read(fd,buffer,len);
    if (result < 0) { if (errno != EINTR) break; }
    else { buffer += result; len -= result; }
}

and

int result;
result = ioctl(fd,cmd,addr);

→

int result;
do { result = ioctl(fd,cmd,addr); }
while ((result == -1) && (errno == EINTR));
```

12.2 Others

12.2.1 I seem to be unable to compile anything with gcc. Why?

If you have only 2 MB RAM, gcc will die silently without compiling anything. You must have at least 4 MB to do compilations

BTW Since swapping is possible, I have heard that compilation works with only 2Meg and a lot disk traffic :) Isn't it great?

12.2.2 gcc complains about not finding crt0.o and the system include files What am I doing wrong ?

The include files normal place is in `/usr/include`. `lib*.a` and `*.o` should be in `/usr/lib` or `/usr/local/lib`

12.2.3 I tried to port a *new* version of gnu stuff. But in the linking phase, gcc complains about the missing libg.a.

Yes this is well known for compiler version earlier than 2.2.2, throw away the flag `-g` that's all, anyway `libg.a` is *only* for debugging purpose.

12.2.4 How can I get bsd style signal?

Use `-D__USE_BSD_SIGNAL`.

12.2.5 How to compile programs which may be debugged with gdb?

There are different ways to handle this problem. If you have the `gcc2.2.2` or later it's simple, use the `-g` flag. Otherwise there are different possibilities:

1. As there is no `libg.a`, you should throw away the `-g` flag in link phase, this means that the compilation must be done in two steps example: instead of `gcc -g monprog.c -o monprog`, use the following `gcc -g -c monprog.c` and then `gcc -o monprog monprog.o` Alas this method is not that good if you are using `Makefile`.
2. The other way is to create an empty `libg.a` as follows (Peter MacDonald's trick):
 - create `libfake.c` containing `libgfake(){}`
 - compile it with: `gcc -c libfake.c`
 - create the `libg.a` with: `ar r libg.a libfake.o`
3. The more tricky Humberto method:

```
cd /usr/lib
ranlib libg.a
```

then `gcc -g monprog.c -o monprog` will produce a debuggable `monprog`.

12.2.6 When compiling some code, cc1 complains about some `insn` code, what's that?

An `insn` is an internal representation that `gcc` uses when compiling. The main part of `gcc` is to take ordinary `c` (or `c++`) code, and compile it, while doing optimizations in `insn` part, which is soft/hard independant. Then another part which is hard/OS dependant takes the `insns` and translate it in assembly language. The fix is only to turn off the optimization flag (`-O`).

12.2.7 When compiling `$#!`, I've got some problems with `SIGBUS` signal that doesn't exist. Any clue?

`SIGBUS` is a common problem, its not needed, just comment it all out, something like:

```
#ifdef SIGBUS
.. normal sigbus code ..
#endif
```

12.2.8 How can I write codes suitable for building shared library?

There are some guidelines for writing codes suitable for building the shared library with jump table.

1. Never, ever allow library users to access global data directly. Always, always make them go through access functions. That way, you are free to change the internal implementation, but can easily provide backward compatibility by simply replacing the access functions.
2. If you do have to define some global data for library users to access, put them into a separate file. Never mix them with other library code. Also make sure the relative locations of the global data will not change very often. One solution is leave some spaces for them to grow.

13 SCSI Special

This section is written by Drew Eckhardt, mail him for any information, questions related to this section.

Last update May 1993.

13.1 General Information

13.1.1 What SCSI hosts are supported?

The Adaptec 154x, Adaptec 174x, Future Domain 8xx (TMC 950 based boards?) and 16x0, Seagate ST0x, Ultrastor 14F (some of the new ALPHA code makes attempts to deal with the 34F), and Western Digital 7000 are supported. Various Adaptec clones from Bustek and Future Domain are known to work, in both ISA and EISA flavors.

There is an alpha driver for the Ultrastor 24F at `tsx-11.mit.edu` in `/pub/alpha/-scsi/u24f-driver.tar.z`.

Some of the SCSI drivers will not autodetect your SCSI host if the BIOS is disabled, and there may be IRQ, DRQ, address restrictions compiled into the distribution kernel.

Virtually all SCSI disks, CD ROMS, and tapes should work.

13.1.2 What SCSI hosts are not supported?

ANSWER: The Adaptec 152x, Always IN-2000, DTC, Mylex, and Ultrastor 24F, PS/2 SCSI boards, and all SCSI ports on sound boards.

If you want to run Linux on this hardware, you'll have to get technical docs on it, and write a driver yourself or bribe someone to do it. Someone is working on a driver for the IN-2000

13.1.3 Which disks?

- Anything that works with your host adapter electronically, and is new enough to support a December 1985 draft of SCSI-I.
- Sysquest, and other removeable hard disks are supported.
- Extended partitions also work, with 11 Extended partitions per disk allowed in combination with the four real partitions and whole disk sub device.
- Disks up to two terabytes in size will work, since the sd driver switches to 10 byte reads when necessary.

13.1.4 What about SCSI tapes?

Tapes are supported. You may wish to obtain the utility program `mt`, which is usually available from `tsx-11.mit.edu` in `/pub/linux/ALPHA/scsi`.

13.1.5 What about CD ROMS?

Yes, CD-ROMs are supported, and the ISO-9660 file system is included with the kernel distribution. Rockridge extensions are supported. You will have to make sure that you have configured the kernel to include the `isofs` filesystem or otherwise you will not be able to use the `cdrom`.

13.1.6 What about flopticals, Bernoulis, Sysquests, and other removable media devices?

The normal SCSI disk driver supports these. Note that pre-.98pl4 versions of Linux had some bugs handling disk-change.

13.1.7 How do I get SCSI information?

Subscribe to the SCSI channel of the `linux-activists` mailing list. mail `linux-activists@joker.cs.hut.fi` And put in the header. **X-MN-Admin: join SCSI.**

13.2 Others

13.2.1 I can't make a filesystem on `/dev/hd*`!

`/dev/hd*` aren't your SCSI disks. `/dev/sd*` are. See section 13.2.7 for appropriate major / minor numbers if they do not exist on your root diskette.

13.2.2 How do I partition the disk?

Use `partdisk` or the DOS partitioning program of your choice.

13.2.3 My partitioning program can't figure out the disk geometry.

The problem with partitioning SCSI disks and Linux is that Linux talks directly to the SCSI interface. Each disk is viewed as the SCSI host sees it : N blocks, numbered from 0 to N-1, all error free. There is no portable way to get disk geometry.

However, DOS doesn't like things like this, it demands that BIOS present it with a normal Cylinder / Head / Sector coordinates. So, BIOS does, and it comes up with some fabrication that fits what DOS wants to see. You don't want to disagree with what BIOS thinks when you write the partition table.

The newest SCSI code will return the mapped geometry for some host adapter / disk combinations. Kernel release 0.99 and later should have this capability.

13.2.4 The linux partitioning programs don't work.

Some of these default to `/dev/hd*`, which are disks on WD-1003 compatible controllers (IDE, MFM, RLL, ESDI, etc), rather than `/dev/sd*` (SCSI disks). Your solutions are to

1. Call the partitioning program with a device name, ie

```
partdisk /dev/sda
```

2. Make links from `/dev/hd*` to `/dev/sd*`

13.2.5 The stock {f,pf}disk programs don't work.

At least some older versions of these partitioning programs have `/dev/hd*` hard-coded into them, and they don't see `/dev/sd*`.

If you want, YOU can change the `/dev/hd*` devices to be links to or other entries for the SCSI devices.

Also, `pfdisk` should be callable with a device name.

Hopefully, some one will change the partitioning programs to respect the SCSI devices.

13.2.6 Where is the latest version maintained?

It is at `tsx-11.mit.edu` in `/pub/linux/ALPHA/scsi`, but join the list before you grab anything.

Eric Youngdale (`eric@tantalus.nrl.navy.mil`) added scatter/gather support to the SCSI drivers, and changed around the internal structures to get something that works as *version 2*.

13.2.7 What are the major / minor numbers for SCSI drives?

Because of the large number of devices that can be hung off of a SCSI bus (as many as 56 if you use SCSI fanouts or bridge boards), and the possibility of 16 partitions on a SCSI disk, we'd run out of minor numbers if they were statically allocated — so a dynamic numbering scheme is used.

Block device major 8 is used for SCSI drives, 11 for CD-ROMs. Character device major 9 is used for SCSI tapes.

Minors are assigned in increments of 16 to SCSI disks as they are found, scanning from host 0, ID 0 to host n, ID 7, excluding the host ID. Most hosts use ID 7 for themselves.

A minor where $\text{minor mod } 16 = 0$ is the whole drive, where $\text{minor mod } 16$ is between 1 and 4, that partition, extended partitions dynamically assigned from 5 to 15 inclusive. Note that the `gendisk.c` module prints partition tables on initialization — you should be able to see them there.

Example : I have four SCSI disks, set up as follows

```
Seagate ST02,  ID = 0
Seagate ST02,  ID = 5
Ultrastor 14,  ID = 0
Adaptec 1542,  ID = 0
```

The first disk on the seagate at ID 0 will become minors 0-15 inclusive, the second at ID5 16-31 inclusive, the disk on the Ultrastor 32-47, on the Adaptec 48-63.

13.2.8 How do I reduce kernel bloat and eliminate the drivers I don't want?

For kernel release 0.99 and later, just go to the `to` directory in the kernel source tree, and type `make config`, and answer the questions. For older kernel distributions, simply `#undef CONFIG_DISTRIBUTION` in `include/linux/config.h`, and define the macros for the SCSI hosts you want enabled.

13.2.9 I get SCSI timeouts.

Make sure your board has interrupts enabled correctly.

13.2.10 The * driver doesn't work

A source of difficulty is that you may be trying to run the bus too fast. Try turning off the turbo switch for your machine, or check the bus speed setting in your CMOS.

13.2.11 My Seagate / Future Domain TMC-88x board is not detected.

The Seagate and Future Domain boards have memory mapped registers. To detect them, Linux scans for a signature in the ROM BIOS (typically, a copyright message) and sets the register addresses relative to that.

This can fail for two reasons

1. If the BIOS is disabled.

In this case, you should edit `kernel/blk_drv/scsi/Makefile` and add `-DOVERRIDE=x -DCONTROLLER=y` where `x` is the base address of your controller (the factory default setting is `0xc8000`) — note that this is *not* the segment (ie, `0xc800`), and `y` is the controller type, either SEAGATE or FD.

2. If we don't know about your BIOS yet

Please use DOS and DEBUG to find us a signature that will detect your board: I.e., if your board lives at `0xc800` do

```
debug
d c800:0
q
```

and send me (drew@cs.colorado.edu) the nearest convenient ASCII message, with the length and offset from `c800:0` or wherever.

13.2.12 The Seagate / Future Domain TMC-88x driver doesn't work.

There are several possibilities

1. Is the board jumpered for IRQ5?

The factory settings are for MSLOSS, and have interrupts disabled. Interrupts are controlled by the W3 (ST01) or JP3 (ST02) jumper, which should have pins FG shorted. See your manual.

2. Cached machines will not have problems IF the Seagate's address space (typically `C8000 - CAFFFF`) is not marked *non cacheable*. This applies to the i486 internal cache as well as i386/i486 external caches.

This can be set in the XCMOS of most machines. If you can't disable cache for the Seagate's area (16K in size, starting at the base address), then you must disable the cache entirely, otherwise it won't work.

3. The new seagate drivers (not yet released) can operate in a mode where the 0ws ISA bus line is used to synchronize the ISA and SCSI busses. This makes things fast, since the alpha Seagate driver can dump SCSI data out fast enough to swamp the bus. However, this will only work if the 0ws jumper is enabled. If this isn't the case, the Seagate driver *should* detect this as an over/under run condition, but it might not. Either enable the 0ws jumper, or `#undef FAST` in `seagate.h`.

13.2.13 When using a Seagate / Future Domain TMC-88x, why does my system hang when syncing to disk?

The Seagate boards are an incredibly brain dead piece of hardware. They can only generate an interrupt when a target raises the SEL signal. So, as long as a target is connected, the Seagate driver must spin its wheels waiting for the actual data transfer. Some devices aggravate the situation by connecting for long periods of time while not doing anything.

13.2.14 The WD-7000FASST driver doesn't work

According to Keith Smith,

There are two different versions of the WD7000/FASST2 One uses a WD33C93A the other a WD33C93. Firmware incompatibilities in the board bios could cause a problem as the former chip fixes some problems that were discovered on the latter. We're talking BOARD firmware, as well as BIOS firmware.

13.2.15 My Ultrastor 14F or 34F isn't detected

The Ultrastor 14F driver won't probe for an Adapter at address 0x310. Either move it do a different address, or recompile the kernel, adding a rule to `kernel/blk_drv/scsi/Makefile`:

```
ultrastor.o: ultrastor.c
$(CC) $(CFLAGS) -DOVERRIDE_PORT=0x310 -c ultrastor.c
```

13.2.16 The Adaptec driver doesn't work.

Older versions of the SCSI drivers required that the Adaptec BIOS be enabled because it does some initialization that affects the driver. If the BIOS was disabled, it may fail on your system. The new scsi code does not require this.

A common source of difficulty is a conflict between two different boards forand IRQ level, a DMA channel, or an I/O address. Check the settings for the boards you have in your system (music boards are zknown to use similar IRQ or I/O addresses as the Adaptec. In the new scsi code the Adaptec can use any of the IRQ levels that it can be strapped for, it can use I/O address 0x330 or 0x334, and it can use DMA channels 5, 6 or 7.

Older versions of the SCSI drivers ran the 17 x in 1542 emulation mode, where some of these cards have a bug that interferes with reads / writes > 512 bytes.

Older versions of the SCSI drivers ran the 174x in 1542 emulation mode, where some of these cards have a bug that interferes with reads / writes > 512 bytes.

The new SCSI drivers support the 174x in *native* mode, so this isn't a problem.

13.2.17 My tape drive or other removeable media device isn't recognized at boot time

Try booting with a tape in the drive.

13.2.18 My system is dog slow (ie, 60k/sec)

SCSI commands have an incredible amount of overhead. For every command, you need to arbitrate for the bus, select the target, establish an L_T_L nexus, and send the command. Processing of that command may take as much as 1ms on older

devices. Add this overhead to what you already have coming through the file system, buffer cache, etc, and you have a real problem.

To work around this, we needed to maximize the amount of data that could be transferred in a single command. So, we implemented scatter-gather, which allows reads/writes from/to contiguous disk sectors to non-contiguous buffers.

This typically gets you a 3-5 fold improvement in performance.

The current kernel only has scatter-gather support for the Adaptec, Western Digital, and Future Domain 16xx boards. However, Alpha code that provides scatter / gather for the Seagate is available from `headrest.woz.colorado.edu` in `/pub/linux/patches.seagate.Z`, for the Ultrastor from `tsx-11.mit.edu` in `/pub/linux/ALPHA/scsi/ultrastor.fast`.

13.3 Bugs

13.3.1 What about bugs?

It works on MY hardware. It works on Tommy Thorn's Adaptec system, Dave Gentzel's Ultrastor, Thomas Wuensche's WD7000, Rick Faith's Future Domain, and quite a few other systems.

Bugs that are there will mostly be very hardware specific, and nasty to track down. SCSI should be basically error free — consequently, the error code has not been heavily tested, and there are known bugs in it.

If you have found a bug, please mail it to the SCSI channel of the mailing list with specifics of your hardware. Other people may have the same problem, a solution, etc.

If you provide a patch, in context diff form, the bug will be fixed Immediately. If you can provide me with a procedure that reproduces the bug on MY system, the bug will be fixed *soon*. Your chances of getting it fixed increase exponentially with the amount of information provided. If I can't reproduce it, and your description gives me no idea as to where in the code to look, it won't get fixed by me.

13.3.2 What are the known bugs?

Many of the drivers cannot abort a command that is in progress — if something goes haywire, all drives off that controller will hang.

13.3.3 I get a message saying READ CAPACITY FAILED. What does this mean?

UPGRADE.

13.3.4 What do I do if I find a bug that still looks like a bug after I've read the FAQ?

Your best bet is to send it to the SCSI channel of the mailing list, where it will be seen by all of the people who've contributed to the SCSI drivers.

In your bug report, please provide as much information as possible regarding your hardware configuration, and all of the messages that Linux prints when it boots. Your chances of getting the bug fixed increase exponentially with the amount of information provided.

The bottom line is that if we can't reproduce your bug, and you can't point at us what's broken, it won't get fixed.

13.3.5 I've found one of the following bugs

- I can't swap to a SCSI disk, or mount one as /
- I get a READ CAPACITY FAILED message on bootup.
- I have a removeable disk (ie Sysquest) and have problems when I change media.
- I have a Seagate / Future Domain TMC-88x and the kernel panics with a kernel paging message.
- I have an Adaptec 1742 and am experiencing data corruption
- I have an Insite floptical drive and it won't work.
- I have a TANDBERG TDC 3600 revision U07, SONYCD-ROM CDU-541 revision 4.3d, DENON DRD-25X revision V, or a SEAGATE ST296 revision 921 and the system hangs or reports multiple devices.

All of these bugs have been fixed, so UPGRADE. The first one was due to a race condition. The fourth was caused by a kernel stack overflow. The Adaptec problem was caused by a firmware bug in 1540 mode. The last two were caused by SCSI devices that weren't compliant.

14 X11 — the Maximum and More

This section is maintained by Krishna Balasubramanian (balasub@cis.ohio-state.edu). Mail him if you have corrections, additions, etc.

Last update May 1993.

The X11 directories are (depending on your version of X)

XLIB = `/usr/X386/lib/X11/` or `/usr/lib/X11/`

XBIN = `/usr/X386/bin/` or `/usr/bin/X11/`

XDOC = `XLIB/etc`

14.1 X386 General Information

14.1.1 What is the X11 release supported by Linux?

It's the X11R5 (xfree86-1.1).

14.1.2 What is X386/xfree86?

X386 is the port of the X11 server to System V/386 that was done by Thomas Roell (roell@informatik.tu-muenchen.de). It supports a wide variety of VGA boards. X386 1.2 is included in MIT's X11R5 distribution.

The Linux X386 port was based on the stock distribution from X11R5, from MIT and was done by Orest Zborowski (obz@sisd.kodak.com).

It has since moved to becoming part of the standard xfree86 distribution.

14.1.3 Where can I get X386 1.2 (X11R5)?

The X386 1.2 and xfree86 sources are available at any site that distributes the X11R5 source (too numerous to list here, but includes export.lcs.mit.edu)

14.1.4 Any tips on compiling X11R5?

- Dont do it.
- XFree86 is distributed with a link kit so you can optionally include what you like in the server.
- You will need lots of disk space (around 150Meg for the MIT core).
- Look at `mit/server/ddx/x386/README\verb` and follow any instructions for your particular platform.
- Contact the X11 channel at `linux-activists`

14.2 Hardware Requirements

Approx: at least 4 megs of ram + swap ...slooooww. 10 Meg disk for X. Another 6-10 meg of disk for GCC if you want to compile X11 programs.

14.2.1 What VGA boards are supported?

- ET3000, ET4000, GVGA, PVGA1a, WD890c00, TVGA8900, ATI ver. 5 or 6, 8514/A, generic VGA: 800x600 Mono (Xmono), 640x480x16 (Xega).
- Diamond cards are not supported and will not be supported. If you are the unfortunate owner of such a card, you can probably get the server up by booting in specific modes or using dos to set your modes before warm booting into linux or using an external clock setting program. You will have to bear with these irks until you can convince diamond to alter their policy.

Standard x11v1.1 or xfree86 server:

ET3000 (for ex. GENOA 5300/5400)
 ET4000 (Tricom, STB PWR Graph, Sigma Legend, etc.)
 GVGA (Genoa 6400)
 PVGA1A (Paradise VGA Professional) ... see BUGS.
 WD90C00 (Paradise VGA 1024)

supported by xfree86:

TVGA TRIDENT 8900c, 9000, support is in xfree86

ATI See the ATI section (??) below.

Those with 8514 compatible cards may want to get the X8514 server for speed (~2x xstones?).

MONO Any vga card should be able to use X386mono server. (At least 640x480 with 800x600 virtual). Use the vga2 section of Xconfig. Support for hercules monochrome card (usable as 2nd display).

Courtesy Rik Faith (faith@cs.unc.edu)

The following servers will usually not handle all the options supported by xfree86. Read the documentation that comes with them carefully. They will be merged with the xfree86 distribution in some time.

8514 ATI graphics ULTRA, ATI graphics Vantage Should work with any VESA standard 8514/A register compatible card? Courtesy Kevin Martin (martin@cs.unc.edu).

Scott Laird (lair@midway.uchicago.edu) writes: I uploaded a new version of the X8514 X Server to sunsite and tsx-11. It is in `/pub/Linux/X11/X-servers/-X8514/X8514scale.tar.Z` on sunsite. It's linked with version 4.2 of the

jumpable libraries, includes TCP/IP support, support for compressed bitmap fonts, Type 1 and Speedo scalable fonts. There's a **README** file in the same directory that will answer more questions.

XS3 S3 chipset server (Jon Tombs (jon@robots.ox.ac.uk)) Get the FAQ on <ftp.robots.ox.ac.uk> (/pub/linux/S3 check sunsite).

Xega Generic 640x480x16 compatible server (originally for laptops). This requires a microsoft mouse at `/dev/mouse` for now and it does not use Xconfig so use environment variables to define the font path etc. in `.xinitrc`:

```
export FONT_PATH=/usr/lib/X11/fonts/misc:/usr/lib/X11/fonts75dpi
```

Works better with courier fonts so add to `.Xresources`:

```
*Font: --courier-medium-r*--10*
```

..or whatever.. A link kit is available at `tsx-11` (you need `gcc2.2.2`). in `/pub/linux/ALPHA/Xega/X386.ega.T.Z` (Obselete? Unusable? Anyone?)

Do not try to bring up an Xserver that does not support your hardware. There have been cases where damage has resulted from pushing the monitor (specially fixed frequency monitors) beyond its specs.

14.2.2 What Mouses are supported?

- Serial : Logitech, Microsoft, MouseSystems compatibles ...
- Busmouses : Logitech, microsoft, ATLXL, PS/2 (aux).

14.2.3 Has anyone gotten the *Mini-DIN* mouse on an HP Vectra 486/33T to work? The slight info I've been able to find says it's PS/2 compatible.

First you need to create an entry in `/dev` for it:

```
mknod /dev/psaux c 10 1
```

There is no direct PS/2 mouse support in XFree86 yet, so to use the mouse with X you'll have to use the `mconv` mouse protocol conversion utility, which can be found on nic.funet.fi, in `/pub/OS/Linux/utils/tools/mconv.c`. This program converts the packets sent by the PS/2 mouse into the corresponding ones from a Microsoft mouse, so you can fool X telling it you have a Microsoft serial mouse instead. Instructions for use are included in the source file. (Johan Myreen, jem@cs.hut.fi)

14.2.4 I have trouble with my logitech Pilot mouse and X under Linux, any clue?

(Thomas Roell?) There are TWO lines of Logitech mice out there. One is the programmable and uses MouseSystems protocol at startup. X386 reprogramms them to use another protocol. If you specify **Logitech** in the `Xconfig`, X386 assumes a mouse like C7 or S9 (not C7-M). The second line is the MicroSoft compatible. Currently all newer Logitech mice follows this practice, like the MouseMan. In that case you have to say **MicroSoft** or **MouseMan**.

Jean Liddle (jliddle@rs6000.cmp.ilstu.edu) just purchased a new logitech mouseman, and yes, the new X-Windows (XFree86 as opposed to X386) requires that you use the "microsoft" mousetype. However, if you turn Third Button Emulation off, the middle button WILL work.

14.3 Linux Distribution

14.3.1 What is the current version?

The newest public Linux release is x11v2.1. It is the Linux release of xfree86-1.1. It was compiled with gcc-2.2.2d and is available for example at `tsx-11` in `/pub/linux/packages/X11/xfree86-1.1/` The SLS distribution stays at 2.0.

14.3.2 What Files do I need to download?

For x11v2.1 there is a README file in the distribution directory. Bring this down first and read it. Then get `xbin`, `xman` and `xfonts`. You need `xprog` and `xlibman` only if you plan to compile X applications. You dont need `xkit` if you want to relink the server.

You need the shared images `libc.so.4` (4.1 or later) and `libm.so.4` (4.0) in `/lib/`. All the X11 binaries use these. To compile programs you should get `gcc-2.2.2d7` or later and the 4.x libraries.

For Other servers see also the notes in the previous section.

14.3.3 Where do the X11 files go? What are they?

The linux X binary distribution looks something like:

```

/usr/X386/
  lib/
  libX11.a ...      (X libraries needed for compiling stuff).
  X11/
    config/        (template files for compiling ... xmkmf)
    fonts/
    misc/ , 75dpi/ ...
    etc/           (documentation).
  bin/            (X386, X386mono, xterm, X linked to X386).
  include/X11/    (include files needed for compiling).
    bitmaps/      (icons, bitmaps used by some applications).
  man/
  man1/, man3/, cat1/, cat3/
/lib
  libX???.so.??   (Shared images needed by all X11 binaries).
```

Its best you do a `tar -tvzf` on the distribution you get or find some means of extracting the distribution file list. The support files include default/example `Xconfig`, `xinitrc`, `twmrc` which you could copy to your `$HOME` directory (as `.Xconfig`, `.xinitrc` and `.twmrc`) and edit them to define your hardware and X11 setup. The directory `XLIB/etc/` contains much documentation on how to configure video modes.

Read the files `XDOC/README` and `XDOC/README.Linux`.

14.4 Learning/Using X

14.4.1 Where can I find the basic help for learning/using X?

Try `man X386`; `man Xserver`; `man xterm`; `man twm`; `man xinit` on any UNIX machine. The man pages on `xterm` and `twm` will reflect behaviour on Linux systems quite accurately. The others will give you some bearing though not every option will work. These man pages are part of the `xfree86` distribution.

14.4.2 What docs are available besides man pages?

The FAQ in `comp.windows.x11` is extensive and has a bibliography. There is also the X-FAQ. Both are available from `export.lcs.mit.edu` in `/pub/contrib/`.

14.4.3 How do I start up X?

Type `startx` at the shell prompt. Before doing this you should:

1. Untar the `xfree86` distribution from `/` (read `XDOC/README`, `XDOC/README.Linux`).
2. copy the file `Xconfig.sample` from the `XLIB/` directory to your `$HOME` directory (`/usr/root/Xconfig?`) or to `XLIB/Xconfig`.
3. Edit `Xconfig` and set the video data for your card and monitor. Use the clocks appropriate for your card and the modes appropriate for your monitor for those clock values. To determine what values are appropriate the `Xconfig` section below and the files in `XDOC/`.
4. Set the mouse device in `Xconfig` correctly....Use the `busmouse` entry for Logitech only. For other busmouses use `microsoft` (try all :-).
5. Ensure that there is a `termcap` entry for `xterm` in `/etc/termcap` (One is available in the subdir `/doc/`).
6. Ensure that the X11 binaries are in the path example: add the line `PATH=${PATH}:-/usr/X386/bin` to `~/.profile` or look at `bash.ad` in `/doc/`.

WARNING: NONE of the modes (in the Modes line in `Xconfig`) should use a clock your monitor cannot handle. This can blow up your expensive monitor!

Now type `startx`.

14.4.4 How do I support national keyboards in X11?

Put a `.Xmodmap` file into the user's home directory specifying which keys to use for what; or else replace the `/usr/lib/X11/xinit/.Xmodmap` with the national keyboard keys customized. Do `man xmodmap` for more details. enclosed an excerpt of mine (french keyboard):

```
keycode 8 =
keycode 9 = Escape
keycode 10 = ampersand 1
keycode 11 = eacute 2 asciitilde
keycode 12 = quotedbl 3 numbersign
keycode 13 = apostrophe 4 braceleft
keycode 14 = parenleft 5 bracketleft
keycode 15 = minus 6 bar
keycode 16 = egrave 7 grave
keycode 17 = underscore 8 backslash
keycode 18 = ccedilla 9 asciicircum
keycode 19 = agrave 0 at
keycode 20 = parenright degree bracketright
keycode 21 = equal plus braceright
keycode 22 = BackSpace
keycode 23 = Tab
keycode 24 = A
.....
```

14.4.5 How do I configure X .. colors, menu, keyboard?

makes me ... dizzy ...

server video, keyboard

Xconfig: man X386, man Xserver, XDOC/VideoModes.doc, man X386keybd,
man xset, man xmodmap

resources global and application preferences.

XLIB/xinit/.Xresources or **~/Xresources:** man xrdp?

twm menus, title bars, colors .. look and feel.

~/twmrc or **XLIB/twm/system.twmrc:** man twm

xinit startup

~/xinitrc or **XLIB/xinit/xinitrc** — man xinit

applications **XLIB/app-defaults/Xxx** (also **~/Xresources**)

man xxx

You can start with files from any X11 setup. Backup the files you change. See the FAQ on `comp.windows.x`.

14.5 Debugging Startup Problems

14.5.1 How do I start up X?

Type **startx** at the shell prompt. Before doing this you should:

1. copy the file **Xconfig.sample** from the XLIB directory to your HOME directory (`/usr/root/Xconfig?`) or to **XLIB/Xconfig**.
2. Edit **Xconfig\verb** and set the video data for you card and monitor. Use the clocks appropriate for your card and the modes appropriate for your monitor for those clock values. To determine what values are appropriate read files in **XLIB/etc** and the **Xconfig** section below.
3. Set the mouse device in **Xconfig** correctly....Use the busmouse entry for Logitech only. For other busmouses use microsoft.
4. Ensure that there is a termcap entry for **xterm** in `/etc/termcap` (One is available in the subdir `/doc`).
5. Ensure that the X11 binaries are in the path example: add the line `PATH=${PATH}:/usr/X386/bin` to `~/profile` or look at `bash.ad` in `/doc`.

WARNING: NONE of the modes (in the Modes line in **Xconfig**) should use a clock your monitor cannot handle. Now type **startx**.

14.5.2 I cannot type in my xterm

`/dev/console` must have `major=4, minor=0`

```
rm -f /dev/console; ln /dev/console /dev/tty0
```

14.5.3 What are some of the common omissions and errors?

1. Add `/usr/bin/X11` (or `/usr/X386/bin`) to your path
2. Add the `xterm termcap` entry to `/etc/termcap`
3. `ln -s X386 X` in the `XBIN` directory `ln -s X386mono X` (if you want to use the mono server)
4. Run X as root first. This avoids some trivial problems
5. If you change the font path entry X386 is finicky about the syntax
6. Make sure your mouse entry in `Xconfig` points to a legitimate serial device (or busmouse)... usually something like: microsoft `/dev/ttys1` serial devices have major # 4 minor 64,65 ... busmouses are major 10. `ls -l in /dev` will show you the major and minor entries.
7. Need read/write/exec access to `/tmp` directory
8. Use `startx 2> x.err` to log your error messages. If the server sticks and you have to reboot you'll probably loose this stuff... see notes on *hanging* below.
9. See Xconfig problems and the device list below.
10. Have you read the X386 man page?

14.5.4 Why is the server unable to find some of the fonts?

First check that the directories listed in the font path exist and have font files in them. Some of the servers are not set up to use compressed fonts. In such cases you have to uncompress the fonts in the directory and run `mkfontdir ..`. Read the man page on `mkfontdir`.

14.5.5 My server hangs Why?

Quite often it isn't a hang. If the server grabs the screen (and enters graphics mode) and then dies, it may return you to bash *without* restoring the screen. In other words, just because you see no output, don't assume it's running/hanging. The way to test this is to see if you can use a bash command to spin your disk. eg.: `sync; ls -l /bin`.

Also — I *suspect* that (at least for me), you can kill the server if you include the `servnum` option in the `Xconfig`.

You can kill the server when you want to by typing ctrl-alt-backspace. Then a couple of ctrl-C 's in case your stuck in `xinit`.

The two line message

X386 version

(protocol Version

is returned from the X386 programme itself, so if you get this (or if the screen blanks) X must be starting.

14.5.6 When I try to use X11 I got Cannot connect to server or process does not exist, any clue?

The cause might be wrong fontpath variable setup in your `Xconfig`; uncomment and set the correct fontpath variable.

14.5.7 Why cant I run more than 4 xterms? How can I have more than 4 pty's?

Set the number in the header `include/linux/tty.h` and recompile the kernel. Also make nodes :

```
mknod /dev/ptypxx 4 minor (where minor = 128 + xx)
mknod /dev/ttypxx 4 minor (minor = 192 + xx)
```

This is for before linux-0.97pl6.... for later versions you only need to make the nodes.

14.5.8 How does X11 start up?

The startup involves the server (`XBIN/X`), and some programs like `startx` and `xinit`.

Typing `startx` runs the script `XBIN/startx`. Look through `startx`. For explanation of the server arguments try `man Xserver`. `startx` does little other than gather arguments and then call `xinit`.

`xinit` calls X (X is linked to X386). `xinit` therefore continues to run, and quite often you will receive error messages from `xinit` rather than X itself. If you get the message `giving up`, it means X has died (quite probably it died immediately) and `xinit` has been unable to start X applications. `xinit` processes the startup script `xinitrc`.

14.5.9 How can I see what startx passes to xinit, and xinit to X?

Add the line: `set -x` near the top of `startx`.

To see what `xinit` passes to X, try:

`rm XBIN/X` — yes, if you have set it up correctly, X should only be a link so it's safe to remove it. Then create a shell script `XBIN/X` containing:

```
#!/bin/sh
echo $*
```

and run `startx`. You should see something like:

```
X :0
```

which is all `xinit` passes to X. If you don't see the `:0`, you may not have your `DISPLAY=:0` set correctly. THEN RESTORE X!! :

```
cd XBIN; rm X; ln -s X386 X
```

14.5.10 What devices does X depend on?

Check your device numbers with `ls -l /dev`. You should have:

```
5 0 tty
4 0 tty0
4 0 console
4 1 tty1 (etc)
1 1 mem
```

```
4 128 ptyp0 pseudo ttys used by xterm/emacs etc
```

```

4 129 ptyp1 (etc) to talk to unix programs.
4 192 ttyp0
4 193 ttyp1 (etc)

4 64 ttys1 one of these is the mouse
4 65 ttys2 or you have a busmouse.
10 x mouse busmouse x = 0 => logitech 1 => ps/2 (aux)
2 => microsoft 3 => ATI_XL.
/tmp/.X11-unix/X0 socket used by X apps to talk to the server.
this is created by X386.

```

14.6 Xconfig and Video Mode Settings

14.6.1 What do the mode names mean?

The mode names like `640x480` are used to match entries in the modes line with the `modeDB` entry which actually defines the mode. You can toggle between modes in the modes line with `ctrl-alt-numericplus`. If you're experimenting with `640x480` you can define modes `640x480a` `641x480b` `642x480c` with corresponding entries under `ModeDB` so you can toggle between these three modes. Try `man X386`.

14.6.2 How do I compute the numbers for the video modes?

See if there is an entry that matches your chip/monitor in `XLIB/etc/modeDB.txt`. Use the clock entries for your chip and the mode entry for that clock appropriate for your monitor. If you cant find a mode entry for a particular clock `needed_clock` but can find a mode entry (for your monitor) for a nearby clock value `found_clock`, scale all numbers in the `modeDB` entry by $(\text{needed_clock} / \text{found_clock})$. `man X386` and read `CONFIG` and `Xconfig.notes`.

14.6.3 How can I determine the clocks on the video chip? Where can I find `clock.exe` or `clock.pas`?

The XFREE Team discourages the use of `clocks.exe` and wants it to be removed from all docs, since the new Server is better at guessing the clocks than `clocks.exe`.

You can get X to tell you what it finds by commenting the clocks entry in `Xconfig`. Use `startx 2> x.err` to log the output. If X starts up this file gets the clocks only after you exit.

I strongly advise using the `modegen` spreadsheets to generate the appropriate settings. If I remember correctly you can find them in `pub/linux/packages/SLS/x4` called `modegen.taz` or something equally as obvious ... Good luck. (Andy, `asb@cs.nott.ac.uk`)

The clocks entry in `Xconfig` serves as an index. As an example:

```

clocks 1 2
ModeDB
1 640x480 .....
2 800x600 .....

```

means the first clock on the chip will be used for the `640x480` mode and so on. Usually one uses the MHz value for the clock for convenience instead of 1 2 you probably have 25 28 ... etc.

14.6.4 What are the settings used with a trident-8900C?

To get X/SLS/Linux working with a trident-8900C, do the following:

1. From Frank Houston (fh8n@uvacs.cs.Virginia.EDU)

```
vga256
Virtual 1024 768
ViewPort 0 0
Modes "1024x768" "800x600" "640x480"
```

Use the VESA monitor timings from the `modeDB` in the example `Xconfig`. You do not need to specify the chipset or clocks.

2. From Henk Vandecasteele (henkv@cs.kuleuven.ac.BE)
TRIDENT TVGA 8900C card with BIOS C2.11, A cheap color monitor AOC (CM-335) with a maximum horizontal frequency of 39.5

```
vga256
Virtual 1024 1024
ViewPort 0 0
Modes "912x684" "800x600" "640x480"
# This mode drives my monitor to the limit (VESA timings for the rest).
"912x684" 45 912 944 1104 1136 684 684 690 712
```

14.6.5 What can I do with a trident 8900B?

The answer comes from Drew Eckhardt (drew@cs.colorado.edu)

1. While the 8900B works with the Xfree 8900C driver, you must specify a chipset line for the 8900C in your `Xconfig`.
2. The trident's are miserably slow in a Color mode. If you don't mind black and white, you should use X386mono. Not only is it faster, but it requires less video memory, and you'll be able to run a $1k \times 1k$ virtual desktop. Again, you must specify the Chipset line to use the banked mono server for virtual resolutions $< \sim 800 \times 600$.

14.6.6 Where can I get Xconfig-files for various hardware?

If you have access to e-mail, e-mail bcr@physics.purdue.edu a message with the subject: `help`

If you do create a new `.Xconfig` file or improve on an old one please mail it to bcr@physics.purdue.edu to save others from having to duplicate your work.

An example `Xconfig` file for the ATI card is available from <ftp.cs.unc.edu>: `/pub/faith/linux/x386`.

14.6.7 What are common problems with Xconfig?

If X isn't firing at all, or even if it clears the screen and crashes (see below) you almost certainly have a fault in `Xconfig`.

1. Note that the server X interprets `Xconfig` (not `xinit`)
2. If you have a fault in `Xconfig` it is possible to crash or hang the server without getting an error message.

3. leave the two paths at the top of `Xconfig` commented out when testing.
4. test with only one graphics mode included — one you are most sure of.
5. If you make a mistake in the mouse section, you will (I understand) get an error message, so if you don't get some mouse error, try elsewhere first. (eg `No such file or directory` indicates you have a fault in the `/dev/ttyxx` line, or `/dev/ttyxx` doesn't exist.
6. X looks in your home directory first to find an `Xconfig`. make certain it's using the one you've been working on!
7. You need double quotes (") around your chipset and modes: eg

```
chipset "et4000"
Modes   "640x480" "640x480a" "800x600"
```

and around the `/dev/ttyxx` eg.:

```
microsoft "/dev/ttys1"
```

8. Check your VGA section and `modeDB` first
9. Ensure `ModeDB` clock speeds match the VGA section. Comment out any unused lines in `ModeDB`, eg.

```
clocks 25 36 00 00 00 00 00 00
```

then you *must* have (and only have) lines under `modeDB` corresponding to speeds 25 and 36. If you have one with (say) 62 uncommented out in the `modeDB` section in place of the 36 line, you will cause the server to die immediately with no error message.

10. A syntax error (such as an unrecognised word) in `Xconfig` will cause a death with no warning.
11. `videoram` must equal the amount of display ram on your chip.

14.6.8 What do I use in my `Xconfig` file to use the bus mouse? Logitech doesn't work with my Logitech busmouse.

Here are the mouse lines from my `Xconfig`:

```
BusMouse "/dev/mouse"
```

You may add a sample rate if you want, but chances are, unless the model number on your mouse matches mine: PC-93-9MD it probably will cause your machine to lock up, as I've gotten several reports saying that.

According to the X386 documentation, non-logitech mice do not support a sample rate.

14.7 X-Applications: Compiling X Programs

14.7.1 Does anyone know where I can find the `libobz.a` library?

This library contained the socket stuff, when X386 V0.1 was released. Meanwhile it has found it's way into the kernel, so remove `-lobz` from the link step in the `Makefile`.

14.7.2 I get `_setjmp` undefined when linking `xv` from `.a` files.

You are trying to link a program compiled with `gcc-2.2.2` and `setjmp` has changed since then. Recompile `xv` from the source using the latest `gcc`.

14.7.3 How do I compile an X application that has an `Imakefile`?

Type `xmkmf` to generate the `Makefile`, then type `make`.

14.7.4 How do I compile a program that has no `Imakefile`?

Use compiler options generated with programs that do have `Imakefiles`. Look in `XLIB/config/` for the linux definitions. In particular you need `-fwritable-strings`.

14.7.5 `gcc` complains the X libraries are not found or links static.

Check out the `-nojump -static` options of GCC ... read the GCC FAQ. You can tell `gcc` what directories to search with `-L`.

14.7.6 When using `Makefiles` generated by `xmkmf` why do I get many error messages?

The old `xmkmf` under Linux doesn't like `#` comments in `Imakefiles`. Therefore you must change the `#`'s to `XCOMM`. But be careful. Don't change any pre-compiler directives. (Matthew A. Lewis (mal111@po.CWRU.Edu))

XII.39) (Ton van Rosmalen wrote):

14.7.7 I recently compiled XFree86 1.2A for Linux and it compiled fine with `gcc-2.3.3` and `libc-4.3.3` iff I skipped the `make depend`.

Rik Faith:

The source for `ed.c` is available from (Linux binaries are in `utilb13.taz`) <ftp.cs.unc.edu:/pub/faith/linux/utills/utilb-1.3.src.tar.Z> This version of `ed.c` has a 4096 character line limit and will work correctly when building XFree86 under Linux.

14.7.8 What are the development tools available for X under Linux?

Some of the X development packages available on Linux: (XS = sunsite.unc.edu:/pub/Linux/X11/) (Contact the people who distribute them in case of problems.)

Libraries :

`Xpm`, `tiff` .. are in `XS/libs/`. If you are compiling a package that uses these libraries it is best to pick up the standard DLLs. hooft@chem.ruu.nl, mitchum.dsouza@mrc-apu.cam.ac.uk (Some older libraries are in `XS/libs/oldlibs`).

Interviews : SLS package

Shared libraries are not yet available so `doc` `sin1MB`, `idraw` `~1.5MB`, `ibuild`?

KHOROS :

This is the second release of my patches (`~15KB XS/xapps`) to the Khoros system for image processing. If you want to compile it, you need Khoros 1.0, patchlevel 5 (source distribution) and up to 100MB free disk space (including package 0.).

Wolfram Gloger (wmglo@dent.med.uni-muenchen.de)

MOTIF : XS

Metro Link, Inc. is pleased to announce the availability of OSF/Motif 1.2.2 for Linux at the introductory price of \$199 for a complete runtime and development system (shared libs, UIL, source for some demos + 1 Orielly book). You can order OSF/Motif 1.2.2 for Linux by calling Metro Link, Inc. at (305) 970-7353 (voice) or (305) 970-7351 (fax) or by mailing us at sales@metrolink.com.

SUIT : XS/devel

Simple User Interface Toolkit.

Included are shared libs and static libs. Compiled with the SLS kit. To use these shared libs you need libX11 3.0 and libc 4.3. Rob Robert-Jan Kooper (rob@is.twi.tudelft.nl), kooper@dutiws.twi.tudelft.nl

tk+tcl : XS/devel

Here is tk3.0 and tcl6.5 sources and libraries. Compiled this with gcc2.3.3 I have also included the wish (X11) binary.

Mitch (mitchum.dsouza@mrc-apu.cam.ac.uk, mitch%markab@mrc-apu.cam.ac.uk)

xview : XS/devel

New `libxview.a` and `libolgx.a` that correct some problems with the ones found in SLS (and maybe in the other xview releases, I don't know).

Kenneth Osterberg ??

14.8 SVGA server for ATI boards

Last Revised May 1993 by Rik Faith (faith@cs.unc.edu).

14.8.1 What's new?

For XFree86 version 1.2 (and later versions), there are two changes to the `extregPlusXLandOrMasks` table:

1. `ER_BO` was changed from 0x28 to 0x31. (In the Linux world, this was distributed as a binary patch long ago.)
2. `ER_BE` was changed from 0x08, but only for cards EARLIER than the 28800-5.

Both of these changes allow more 28800-4 cards to work, although either one may be sufficient for any one card. You can find these changes in the driver by searching for the Feb 3 timestamps.

14.8.2 What is the ATI SVGA driver?

The ATI SVGA driver is a 256 color driver for the XFree86 server. The driver was written for the ATI WONDER series of graphics adapters. In general, the following modes are supported for cards with 512kb of memory:

640x480x256

600x800x256

The following modes are supported for cards with 1024kb of memory:

1024x768x256

14.8.3 What is the ATI SVGA driver NOT?

The ATI SVGA driver is **NOT** a black and white (2 bit) driver. Therefore, it will **NOT** work with the XFree86 monoserver.

The ATI SVGA driver is **NOT** a 16 color driver. Even though your manual says that your graphics adapter has a 1024x768x16 mode, the ATI SVGA driver will **NOT** use this mode. If you want to use 1024x768, you **MUST** have 1MB of memory on your board and use the 1024x768x256 mode.

The ATI SVGA driver will **NOT** support more than 8 bits of pixel depth. Therefore, even though your manual says that your graphics adapter supports modes using more than 256 colors, the ATI SVGA driver will **NOT** use these modes.

The ATI SVGA driver is **NOT** an accelerated driver. If you have an ATI GRAPHICS ULTRA series card, the ATI SVGA driver will **NOT** use the accelerated hardware. It will only use the SVGA hardware. This will probably not be any faster than a VGA WONDER card, unless the card is on a local bus (and then, it won't be faster than a VGA WONDER card on a local bus, if a card of this type existed).

14.8.4 What cards will the driver work with?

This is a difficult question, because there are many different ATI video cards, and it is sometimes difficult to tell them apart. The best way is to look on the card itself for the ATI chip numbers. Here is a partial list of ATI VGA WONDER cards. The MOST IMPORTANT number is the CHIP VERSION.

Adapter version	Dot clock	Chip version	
VGA WONDER V3		ATI18800	NOT SUPPORTED
VGA WONDER V4		ATI18800-1	NOT SUPPORTED
VGA WONDER V5	ATI18810	ATI18800-1	NOT SUPPORTED
VGA WONDER + (V6)	ATI18810	ATI28800-2	? MAY WORK ?
VGA WONDER (1 MB)	ATI18810	ATI28800-4	SUPPORTED
VGA WONDER (1 MB)	ATI18810	ATI28800-5	SUPPORTED
VGA WONDER XL	ATI18810	ATI28800-5	SUPPORTED

Note that the ATI SVGA driver code has hooks and some support for all of the cards listed. However, with the earlier cards, there are serious differences in the programmer's interface. The only people who have ever worked on the code have had 28800-4 and 28800-5 chips, so these chips are the only ones that are actually known to work.

The ATI GRAPHICS ULTRA chips seem to have SVGA support which is equivalent to the 28800-4, 28800-5, and 28800-a, although I can find no documentation on this in the technical manuals provided by ATI. The ULTRA PLUS and ULTRA PRO boards have been tested to some extent, and seem to work fine, but the original ULTRA may have some problems at greater than 640x480 resolution.

14.8.5 What should I put in my Xconfig file?

THE MOST IMPORTANT LINE IS THE CLOCKS LINE:

```
Clocks      18 22 25 28 36 44 50 56
             30 32 37 39 40 0  75 65
```

The first row of clocks is usable on all VGA Wonder cards. The second row is usable only on V5, PLUS, and XL cards (28800-2, 28800-4, and 28800-5). However, you should include **BOTH** rows for **ALL** cards.

The specification is different from the specification used with the ET-4000 cards. Programs with claim to detect the clocks usually work only with ET-4000 based

cards. No such program exists for the ATI cards, since it isn't needed. THE CLOCKS LINE SHOULD NEVER BE CHANGED. The order of numbers in the clocks line *is* significant.

The actual frequencies, which can be used to calculate the ModeDB lines, are as follows:

```

18 = 18.000
22 = 22.450
25 = 25.175
28 = 28.320
36 = 36.000
44 = 44.900
50 = 50.350
56 = 56.640

30 = 30.240
32 = 32.000
37 = 37.500
39 = 39.000
40 = 40.000
 0 = 56.640
75 = 75.000
65 = 65.000

```

The Chipset will be automatically detected. The chipset name for this driver is `ati`.

Here is a sample Xconfig (which I use):

```

RGBPath "/usr/lib/X11/rgb"
FontPath "/usr/lib/X11/fonts/misc/,/usr/lib/X11/fonts/75dpi/"

Keyboard
    AutoRepeat 500 5
    ServerNumLock

#Microsoft "/dev/mouse"
#MouseSystems "/dev/mouse"
#MMSeries "/dev/mouse"
#Logitech "/dev/mouse"
#MouseMan "/dev/mouse"
Busmouse "/dev/mouse"
# BaudRate 1200
  SampleRate 150
# Emulate3Buttons

vga256
    ViewPort      0 0
    #Modes "1024x768v" "800x600v"
    Modes        "800x600v"
    Clocks       18 22 25 28 36 44 50 56
                30 32 37 39 40 0 75 65
    Displaysize  260 184

```

ModeDB

```
# From NEC MultiSync 3FGx manual (copied by faith@cs.unc.edu)
# active_area front_porch pulse_width back_porch for horiz and vert
# VESA 640x480 at 72Hz (37.860kHz)
# 20.317us 0.762us 1.270us 4.603us; 12.678ms 0.238ms 0.079ms 0.740ms
# (I couldn't find a good clock rate for this one)
# VESA 800x600 at 72Hz (48.077kHz)
# 16.000us 1.120us 2.400us 1.280us; 12.480ms 0.770ms 0.125ms 0.478ms
"800x600v" 50      800 856 976 1040    600 637 643 693
# VESA 1024x768 at 60Hz (48.363kHz)
# 15.754us 0.369us 2.092us 2.462us; 15.880ms 0.062ms 0.124ms 0.600ms
"1024x768v" 65     1024 1048 1184 1344    768 771 777 806
```

14.8.6 What is the history of the driver?

The complete history of the driver is cloudy. I have tried to reconstruct the history from several email messages from a few of the people involved. The history is probably incomplete and inaccurate.

Apparently, Per Lindqvist (pgd@compuram.bbt.se) first got an ATI driver working with an early ATI card under X386 1.1a. (This original driver may have actually been based on a non-functional ATI driver written by Roell.) Then Doug Evans (dje@cygnus.com) ported the driver to the ATI XL, trying in the process to make the driver work with all other ATI cards.

I (Rik Faith) obtained the X11R4 driver from Doug Evans in the summer of 1992, and ported the code to the X386 part of X11R5. This subsequently became part of XFree86. Per and Doug did the majority of work, making the port somewhat trivial. I am grateful for their contribution.

The port to X11R5 required writing the ATIProbe and ATIEnterLeave routines, and re-organizing the ATISave and ATIRestore routines. Someone else added the operating system specific code, but I don't know who.

The major reference that I used was manual from ATI:

```
"VGA WONDER Programmer's Reference"
ATI Technologies, 1991.
Release 1.2 - Reference #PRG28800
(Part No. 10709B0412)
```

However, Chapter 11 (ATI 18800 ATI VGAWONDER) of George Suttly and Steve Blair's *Advanced Programmer's Guide to SuperVGAs* (Brady/Simon & Schuster, 1990) was also useful. Further, someone e-mailed me a random document from the ATI BBS which was dated 3Jul91 and which may have been named `PROGINFO.DOC`.

14.8.7 What is the future of the driver?

I no longer use an ATI Wonder card and have no interest in this driver. Since I have received no bug reports or patches for the driver during the past 3 months, I think it is about as stable as it is going to get. (I suspect that support for the 28800-2, 28800-4, and 28800-5 chips can be made more robust, but I doubt that the other chips will ever be fully supported.)

I have absolutely **NO** intention of porting the code to the mono or 16 color servers. If *you* want ATI support in the mono or 16 color servers, they *you* should do the port. (Since I no longer use my ATI Wonder card, please feel free to take over the current driver as well.)

14.9 Bugs

- Restoration of the text screen fails on some hardware.
You can get the program `runx` from the `vgalib` distribution (`tsx-11/pub/linux/sources/libs/`). Report the problem to the `xfree86` team.
- Cant compile Xaw programs without `-static` in versions `x11v1.1` or older.
There is a glitch with the Xaw shared libs ... I haven't seen this in a while anyway its fixed in newer releases.
- Color problems with Xega fixed in latest version on `tsx-11`.
- It is possible to run X386 on a Paradise VGA Plus. The X386.0 doc file has a small misprint. The Chipset is not `pvga`. It's `pvga1`. Anyway, under the line in `/usr/X386/lib/X11/Xconfig` `VGA256`, add, `chipset "pvga1"`. You will need a 512k board. There seems to be a problem with 800x600 mode that causes pixels to be misplaced in the center 1/3 of the screen. 640x480 works fine. Anybody out there might have a fix for that? You may want to use `X386mono` with 800x600 instead.
- Fix for missing numlock control in `xfree86-1.1`. use `xmodmap` to change the mapping:

```
clear mod1
add mod1 = Alt_L
keysym Alt_R = Mode_switch
add mod5 = Mode_switch
```

David (Dawes)

Sources for X11-FAQ:

Steve Kotsopoulos, Peter Hawkins, John Morris, MM. Corsini, K. Balasubramanian.
Direct comments, questions, complaints to Krishna at: balasub@cis.ohio-state.edu

15 Networking and Linux

This section should be maintained by Philip Copeland (p_copela@csd.uwe.ac.uk).

15.1 General

15.1.1 Where can I find useful information about networking for Linux?

Join the NET channel of the mailing list. Also, you need to read the NET-FAQ by Phil Copeland, it's on sunsite and `tsx-11` and posted to `c.o.l` and the NET channel every couple of weeks.

Basically, all of the information on networking and TCP/IP for Linux is in the NET-FAQ, because it's maintained seperately than this document (and we don't want to waste the space reprinting it all here — it's big!).

The NET-FAQ has been completely rewritten as of Feb 1993. Get the new version.

16 Everything for printing

This is the `lpd` FAQ version 1.2, this section is maintained by Brian McCauley. Please send any suggestion to B.A.McCauley@bham.ac.uk.

Last update June 1993.

16.1 General Information

16.1.1 What is the lpr/lpd package?

`lpr` is a utility that is used to print files out to local and remote printers. As with all networking, there are configuration files which have to be tailored to your needs. These files are concerned with access control, accounting and page setup of your printer(s). Now all control of the printer is performed by a process called `lpd`

`lpd` is a laser printer daemon (yes I know an epson isn't a laser printer but thats besides the point). The program `lpd` is normally started in `/etc/rc.local` in either the foreground or the background. What happens is that `lpd` will sink back as a daemon, only ever being invoked when a printer request is generated. `lpd` itself has some options which are of possible interest to some people chiefly that of the `-l` flag which causes `lpd` to log requests received from the network (useful for debugging). The other option is for what port `lpd` should listen for data on and is generally not used.

Note:

`lp` really stands for line printer but as technology has advanced laser printer makes more sense.

16.1.2 What is the Linux lpr/lpd?

Up until 28th May 1993, as far as there is a definitive Linux version of anything *the* Linux `lpr/lpd` is Ross Biro's (`bir7@leland.stanford.edu`) port of the BSD sources dated 20th November 1992 (so if you're used to AT&T's you'll find some differences). Ross's binaries and the diffs are found at `tsx-11.mit.edu` in `/pub/linux/binaries/usr.bin/lpr.tar.Z` and on `sunsite.unc.edu` in `pub/Linux/system/Daemons/lpr.tar.Z`

The SLS file `lpr.tgz` on disk b2 contains the same stuff.

However this is all about to change so from now on this version will be called the *old* binaries. As of 28th May 1993 there appears to be two different new releases of `lpd` expected RSN.

In addition to the standard BSD `lpr` capabilities Ross Biro's port should have had better support for serial printers (but it didn't work in the old binaries).

Ross's port was not based on the latest BSD sources (by about 3 years, '88 vs '91). His new release will also be based on these sources which can be found at:

`gatekeeper.dec.com` and `/afs/ir.stanford.edu/src/...`

Recent BSD sources are all over the place.

`ftp.uu.net: /systems/unix/bsd-sources/usr.sbin/lpr`
`src.doc.ic.ac.uk: /unix/bsd-sources/usr.sbin/lpr`

to name but two. The BSD sources have unGNU things in their Makefiles so they don't compile *out of the box* but have now been ported by `waltje@Walt.NL.Mugnet.ORG` Fred van Kempen and will be available RSN (if not now) as part of the complete overhaul of the Linux Net stuff.

Ross agrees that Fred's version should be thought of as *the* Linux `lpd` package. I (`B.A.McCauley@bham.ac.uk`) have not had a chance to look at Fred's port yet (I'm waiting for the full release of 0.99pl10 before I update anything) but most of this FAQ refers to the BSD `lpd` package or Linux in general so this is probably not too important.

There are some other versions of `lpd` stuff about but I have to draw the line somewhere. If anyone wants to maintain FAQ sections on these they are welcome to do so. The other versions include `plp` and the `ka9q lp` facility. There are also rumours of SysV spoolers being ported.

16.1.3 What are all these directories for? (/etc, /usr/etc, /etc/inet, /usr/spool, /var/spool)

This is one to which I can't find a straight answer.

For various reasons including the difficulties of fsc'ing / some people like to put all things that aren't needed during bootstrap in /usr so that / is as small as possible. Also some people like all their internet stuff in /etc/inet. Most Linux gurus advise keeping it all in /etc.

If you really want this quirky layout:

```
cd /etc
ls -s ../usr/etc/* .
ls -s inet/* .
```

Make sure the files in /usr/etc and /etc/inet are not already symlinks into /etc.

If you want to spread out your configuration files for the sake of tidiness then a better solution is to have a directory tree called /conf and make the files in /etc symlinks into this tree.

To avoid the problem of your /usr filesystem filling up as a result of stuff in spool queues some people prefer to keep their spool queues and such like in a separate filesystem called /var. If you don't have a /var file system then you should make /var a symlink to /usr. If you *do* have a /var filesystem then you may want to make /usr/spool a symlink to ../var/spool so that any programs expecting to find /usr/spool will not be disappointed.

16.2 Installing

16.2.1 How do I get lpr to work properly other than for root? Where do the files go and what should their permissions be?

Put /etc/lpd in your /etc/rc or /etc/rc.local (usually in /etc/rc.local after you start syslogd (if you use syslogd)).

Set the group fields of the file permissions/ownership as follows:

```
-rwxr-s--- 1 root daemon 37892 Nov 19 23:32 /etc/lpd
-rwx--s--x 1 root daemon 21508 Nov 19 23:32 /usr/bin/lpc
-rws--s--x 1 root daemon 17412 Nov 19 23:32 /usr/bin/lpq
-rwx--s--x 1 root daemon 17412 Nov 19 23:32 /usr/bin/lpr
-rwx--s--x 1 root daemon 17412 Nov 19 23:32 /usr/bin/lprm
```

...and for each of the spool directories listed in the sd fields of /etc/printcap... /usr/spool/lp1:

```
total 5
drwxrwxr-x 2 root daemon 1024 May 18 23:00 .
drwxr-xr-x 11 root root 1024 Feb 19 20:56 ..
-rw-rw-r-- 1 root daemon 4 May 18 23:00 .seq
-rw-rw-r-- 1 root daemon 18 May 18 23:00 lock
-rw-rw-r-- 1 root daemon 25 May 18 23:00 status
```

Note these 3 files are created by lpr and lpd so if you've never run these they could be missing so touch them into being.

Some of these permissions may be overkill so don't be surprised if your system works with different permissions. The fact that lpq needs to be suid root is inelegant and may change in the future.

Alternative approach: just make `lpc`, `lpr` and `lprm` `setuid(root)` then you can forget the file permissions on the spool queues!

You're free to choose different directories for the executables on your system (notably `lpc` is often in `/etc` even though it has commands that are useful to non-root). The master `lpd` lock file which is always in `/usr/spool` so you must have one of those but you needn't necessarily keep your spool queues there.

The main configuration file is `/etc/printcap` although the old binaries look for it in `/usr/etc/lpd`. The README says `hosts.equiv` and `hosts.lpd` files are looked for in `/etc/inet` but again the more accepted convention is to locate them in `/etc`.

For an simple life:

```
mv -i /usr/etc/inet/* /etc/inet/* /usr/etc/* /etc/
rmdir /usr/etc/inet /etc/inet /usr/etc
ln -s ../etc /usr/etc
ls -s . /etc/inet
```

But make sure the files in `/usr/etc` and `/etc/inet` are not already symlinks into `/etc` or you'll be in deep water.

Eventually you'll be able to delete these symlinks once you've got rid off all the old stuff.

16.2.2 Which `/dev/lp*` is my printer?

On an XT bus system LPT1: becomes `/dev/lp0` (major=6, minor=0), on an AT LPT1: becomes `/dev/lp1` (major=6, minor=1). To be more precise:

Name	Major	Minor	I/O address
lp0	6	0	0x3bc
lp1	6	1	0x378
lp2	6	2	0x278

16.2.3 When the kernel boots it says using polling driver, is there an interrupt driven driver?

Yes, the latest kernels support parallel port IRQs but to enable them you must use `lpcnt1`. The interrupt driver uses less processor time and gives a faster throughput but sometimes one may choose sacrifice this to free up an IRQ line. If your printer is `/dev/lp1` using IRQ 7 put:

```
/etc/lpcnt1 /dev/lp1 7
```

in your `/etc/rc`.

If you don't have `lpcnt1` get it from: `tsx-11.mit.edu: /pub/Linux/BETA/lp/-lpirq.4.tar.Z`

16.2.4 What's the difference between the polling driver and an the interrupt driven one?

A polling sits in a loop repeatedly reading the port hardware to see if it is ready for another character and then gives it one. An interrupt driver goes to sleep and is woken by the port hardware (using a hardware interrupt) when it wants another character. Polling drivers are easier to write but in multitasking environments they are a bad thing.)

16.3 The Printcap Database

16.3.1 Where do I get a printcap for a xxxxx?

This question is essentially meaningless - see next question.

16.3.2 What is the content of `/verb"/etc/printcap`?

Given the similarity in appearance and name between `/etc/termcap` and `/etc/printcap` one could be forgiven for assuming that they contain analogous information. This is not the case. Whereas `termcap` contains information about terminal *types* — (mostly escape sequences) `printcap` contains information about *specific* printers (like the directory that holds the spool queue and the device name of the printer). The information about a printer model's escape sequences and so on are held in the various filters which are *programs* called by `lpd` to drive the printer. `/etc/printcap` simply gives the locations of these filters. Alternatively the Net-FAQ has a summary of some of the more important files. One last point you should always specify *suppress header* :`sh`: unless you have *text* printer and want banners. (See the banners question (??) for further expansion).

16.3.3 What is the format of a `/etc/printcap`?

Ideally look out for the format in `/etc/termcap` (yes, I said em `termcap`) but since most people don't have docs about `termcap` here are the essentials.

Lines starting `#` are comments (as you might have guessed).

For each printer usable from the `lpr` command on your system there is one logical line in the file. For the sake of readability each logical line may be spread over several physical lines by making the last character on all but the last physical line a backslash.

Each logical line has the following format:

```
name1|name2|name3:string_capability=string:\
      :numeric_capability#number:boolean_capability:
```

The leading spaces and colon on the second line are for readability only.

A printer can have as many names as you like but conventionally the final name is used as a longhand description of the printer. (Still people are free to say `lpr -P "Fred's grotty teletype"` if that's the description you've given.)

The list of capabilities can be as long as needed and the order is not significant. Each *capability* is denoted by a two character code. (The name *capability* comes from the file format's `termcap` heritage — parameter or attribute would be a more sensible terms.) [Note from Ross Biro: capabilities with 3 character names don't work properly which is why the serial port stuff in the old binaries failed.] Capabilities having string value and have a `=` delimiter between the capability name and the value while those having a numeric value use a `#` (actually they can use either a `#` or an `=`). Boolean *capabilities* are true if they appear in the list and false if they do not.

Special characters in a string value can be expressed using `\`-escape sequences as in C; in addition, `\E` stands for ESC. `^` is also a kind of escape character; `^` followed by CHAR stands for the control-equivalent of CHAR. Thus, `^a` stands for the character control-a, just like `\001`. `.`, `\` and `^` themselves can be represented as `\.`, `\\` and `^^`.

Example:

```
lp|bam|BAM's Epson FX-80:lp=/dev/lp1:sd=/usr/spool/lp1:sh:mx#0:\
:df=/usr/local/lib/magic-filter/lp.df:\
:if=/usr/local/lib/magic-filter/lp.if:
```

Name is `lp` (this is the printer that `lpr` uses by default). It's also known as `bam` or "BAM's Epson FX-80".

The printer is on `/dev/lp1` (aka AT-bus LPT1:). I don't want a burst page. I don't want a file length limit. Files queued by `lpr -d` are passed through `/usr/local/lib/magic-filter/lp.df` and those queued by `lpr` through `/usr/local/lib/magic-filter/lp`.

See also the next question.

16.3.4 My `/etc/printcap` looks identical to someone else's but it doesn't work — why?

See if `lpc stat` reports a printer called `'.'`. The last character on a continued line must be a `\`. If there are whitespace characters after the `\` then it doesn't register the next line as a continuation.

16.3.5 What's the minimum `/etc/printcap`?

This is a silly question but it *is* frequently asked. The answer is `lp:sh` (that's 6 bytes including the linefeed character on the end). To use this `/etc/printcap` you must make `/dev/lp` a symlink to your printer and create your spool queue directory as `/usr/spool/lpd`. (You might think that if you wanted banner pages you could loose the `:sh` but the `termcap` syntax requires at least one characteristic per entry).

16.4 Common Problems

16.4.1 Why do I get jobs queued, but cannot start `daemon` or `lpc: connect: No such file or directory`?

Firstly this `start` in the context of this error really means `wake`. If `ps ax` does not reveal a `lpd` then you `daemon` has died (or was never started).

If you so have an `lpd` then you are having trouble with the socket connection to `lpd`. I don't really understand this but the problem was common in early days with the old binaries but developments in the TCP/IP code in the kernel fixed it for most people. Using the new `lpd` from the `net-2` package without the `net-2` kernel stuff sometimes gives this problem.

When I get it I find that kill and restarting `lpd` a few times will eventually get rid of the problem. There has also been a suggestion that recompiling your kernel without TCP/IP at all can help. If `linux-0.99pl10` with it's all new TCP/IP stuff doesn't fix this then further reserch is probably called for.

16.4.2 How do I prevent the *staircase effect*?

Unix terminates each line of a file with a linefeed but not a carriage return so taken literally a Unix text file printed on an ASCII device will start each line below the end of the previous line. Some printers can be set to treat "linefeed" as "carriage return, linefeed", others can't. If yours can then do simply do that. If the printer cannot be fixed create a shell script filter that reads:

```
#!/bin/sh
if [ $1 = -c ]; then
    cat
else
    sed -e s/$/\^M/
fi
echo -ne \\f
```

where $\^M$ is a carriage return character not a $\^$ followed by a **M**. To type $\^M$ in emacs use the sequence C-q C-m and in vi use C-v C-m. Conventionally this script is called `/usr/lib/lpf`. The test of `$1` allows the insertion of carriage returns to be switched off by the `-l` switch on `lpr`.

Alternatively your printer may have an escape sequence that will set the way it handles linefeed characters. A simple filter that uses an `echo -ne` command to send this sequence may be appropriate.

```
#!/bin/sh
# Filter for HP printers to treat LF as CRLF
echo -ne \\033\&k2G
cat
echo -ne \\f
```

Ross Biro's new release will include a filter called `lpof` that does this.

16.4.3 Why do files with a lot of graphics get truncated?

Usually because you've got a limit set on the maximum size file that can sit in the spool queue. Put `mx#0` in your `printcap`.

16.4.4 Why doesn't `lpr -i` work?

To get `lpr -i` to work you need a filter installed as `if` that implements it. The `-i` switch is simply passed on by `lpd` to the filter. There is such a filter called `/usr/lib/lpf` that can be obtained from `bsd` source archives in the directory `bsd-source/lpr/filters`.

16.4.5 Why doesn't `lpr -p` work?

Because its broken in the '88 sources from which the old binaries were compiled. `lpd` always thinks that the printer is 0 characters wide regardless of what `printcap` says. This is still broken in Ross's new release and Fred's port of the '91 sources. Looking at the sources the fix for Fred's port appears to be:

```
*** printjob.c.old Sat Feb 27 14:07:01 1993
--- printjob.c Sun Jun 6 11:02:36 1993
*****
*** 265,271 ****
    */
    for (i = 0; i < 4; i++)
        strcpy(fonts[i], ifonts[i]);
!   strcpy(width+2, "0");
    strcpy(indent+2, "0");

    /*
--- 265,271 ----
    */
    for (i = 0; i < 4; i++)
        strcpy(fonts[i], ifonts[i]);
!   sprintf(&width[2], "%d", PW);
    strcpy(indent+2, "0");

    /*
*****
```

```

*** 1081,1087 ****
    FF = DEFFF;
    if ((PW = pgetnum("pw")) < 0)
        PW = DEFWIDTH;
-   sprintf(&width[2], "%d", PW);
    if ((PL = pgetnum("pl")) < 0)
        PL = DEFLENGTH;
    sprintf(&length[2], "%d", PL);
--- 1081,1086 ----

```

This patch is untested as I don't have a complete set of net-2 include files on my machine yet.

One other thing: `lpd` calls `pr` by full pathname so if you keep `pr` somewhere different from where `lpd` looks you will need a symlink. Where `lpd` expects to find `pr` varies from version to version but I think `/usr/bin/pr` is the usual place.

16.4.6 Why does `lprm` say too many users?

This was a bug in the C library in the `scandir()` function which failed when an unsorted directory list was requested. This has been fixed now so just get a new `libc.so`. Alternatively there is a `lprm` with a workaround (by toy@soho.crd.ge.com) at sunsite.unc.edu/pub/Linux/system/Daemons/lprm-fix.tar.z

16.4.7 Why does `lpc` complain that it hasn't stopped a daemon?

There's one `lpd` process that runs all the time and it spawns off children to handle each printer as needed. These children leave lock files lying about so `lpc` thinks that they are still running and tries to kill them. This is probably a bug but it's not serious.

16.5 Special

16.5.1 How do I get my printer to go back to the default font after each printout?

Define the `tr capability` in `/etc/printcap` to be your printer's font reset command. For details of the format of this string see the question on the format of `printcap`. This may not work if a printout crashes in the middle of an escape sequence — putting a lot of `^@` on the front may help but this probably won't be enough if you were printing raster graphics when the filter died.

16.5.2 How do I prevent a formfeed at the end of every printout?

If you don't have an `if` specified in `/etc/printcap` then `lpd` will automatically put a formfeed at the end of each file. If you're using a filter then it's up to the filter to decide if it wants to put a formfeed. To disable formfeed completely if you don't have an `if` put `:ff=:` in your `/etc/printcap`. But please note this suppresses the formfeed that would usually be printed if a filter dies. If you want formfeeds after text printouts but not on printouts printed with `lpr -l` then create the following `if` filter:

```

#!/bin/sh
cat
if [ "$1" != -c ]; then
    echo -ne \\f
fi

```

If you want a formfeed after `lpr -l` to be optional you can misuse the `-i` switch to suppress the formfeed with the following trick (after all `lpr -i -l` would usually not be implemented).

```
#!/bin/sh
cat
# use lpr -i -l to print raw without trailing formfeed
if [ "$1" != -c -o "$4" = -i0 ]; then
    echo -ne \\f
fi
```

16.5.3 How do I get burst/banner pages?

For a simple text printer (in particular not postscript) and a simple text banner simply take `:sh:` out of the `printcap` record. If you want to prevent the banner coming out in whatever font was last used on the printer then define the `tr capability` to be your printer's font reset command.

If you want a fancy customised banner (or have a postscript printer) leave `:sh:` in the `printcap` and make each of your filters print the banner. All the information to put on the banner is included in the filter's positional parameters. RTFM(`printcap`) for details. If you're using my `*B.A.McCauley@bham.ac.uk` `magic-filter` package then call the code to print the banners from the config script.

16.5.4 How do I print text on a ps printer?

You need a filter based on a program called `enscript`. I don't know where you can get `enscript` for Linux.

16.5.5 How do I print over a Network?

To print on the printer `foo` connected to the machine `bar.baz.net` from the machine `mine.baz.net` you put an entry like this in your `/etc/printcap` (on `mine.baz.net`):

```
foo:lp=:rm=bar.baz.net:rp=foo:sd=/usr/lpd/spool/foo:
```

and, of course, create the spool directory `/usr/lpd/spool/foo`.

There's no point specifying filters and the line in `mine.baz.net: /etc/printcap` as it's the one in `bar.baz.net: /etc/printcap` that will get used.

On the machine `bar.baz.foo`, you need to put `mine.baz.net` on a line by itself in either `/etc/hosts.equiv` or `/etc/hosts.lpd`; note that putting it in `/etc/hosts.equiv` will allow for unauthenticated logins as well as printing. `/etc/hosts.lpd` is printing only.

Editors note: I realise there are more questions about net printing but I don't know them. If you find any let me know (with answers if possible). See also the Net-FAQ.

16.5.6 How do I write filters?

In normal Unix terminology, filters are just programs (so they must have execute permission) that read a stream from their standard input and write to their standard output.

`lpd` filters are filters in the sense that they read from their `stdin` and write to their `stdout` but are not necessarily true filters in that they may assume that their standard input is a file and perform `lseek()` operations on it. I'm not sure that such filters are strictly conforming but they seem to work OK.

All `lpd` filters must conform to a particular command line syntax (or more often simply ignore command line parameters). For details of the command line parameters RTFM(printcap).

It's easier to debug filters if you test them in an immediate shell before you install them. (If your filter make use of its command line arguments you'll have to specify them too).

```
my-new-filter <file >/dev/lp1
```

If you want to write a shell script filter it must have a `#!/bin/sh` header. Here for example is my Epson FX-80 dvi filter:

```
#!/bin/sh
/usr/TeX/bin/dvips -f | \
/usr/bin/gs -q -sDEVICE=eps9high -r120x216 -dNOPAUSE -sOutputFile=- -
```

16.5.7 I've written a filter in shell script which works OK when I use it from the shell prompt but fails when I tell `lpd` to use it?

You may need a `#!/bin/sh` header. You may also need to set `PATH` within the script since the daemon's `PATH` may not have everything you need.

16.5.8 When should I define an of filter?

Never. (Well strictly speaking there are circumstances but you're unlikely to meet them until you're so familiar with `lpd` that you won't need this FAQ.)

16.5.9 Where do I get filters for given printers?

From BSD source archives (I think).

Editors note: I need a better answer here.

If you already have a program to print say DVI on your printer then making it into a filter is usually a matter of writing trivial shell script — see this FAQ. If the program you are using insists on reading a names file as input see the next question. Text mode filters are trivial too (see this FAQ) unless you want `lpr` to have a choice of fonts in which case they are slightly harder than trivial. You will probably want to insert and `echo -ne` command at the beginning and end of your filter to set up the font etc. to your liking.

16.5.10 How do I get a program that won't read from `stdin` to act as a filter?

For example `dvilj2p` insists on a named file as its input (and what's more expects one with a `.dvi` suffix). To make a dvi filter for a laserjet you'll have to capture the `stdin` into a temporary file.

```
#!/bin/sh
cat >/tmp/$$ .dvi
dvilj2p /tmp/$$
rm /tmp/$$ .dvi
```

The problem with this kludge is that it may have a very high disk overhead. (In the long term `dvilj2p` should be fixed).

16.5.11 The set of filters supported by lpd seem strange — why aren't there filters for Postscript, GIF, TIFF and so on?

History. You can, in fact, use any of the filters (with the possible exception of the default filter) for any reason. If you're never going to use Benson Varian raster files you could use the `-v` switch for GIF files. If you are on a network remember that the filter setups go on the print server so you will need to consider other people's needs too. To avoid running out of possible types use magic filters.

16.5.12 What are magic filters?

Magic filters deduce their input files' types from *magic numbers* (distinctive byte patterns at particular offsets). Magic filters are usually perl scripts, shell scripts or C programs that simply identify the file type then call the appropriate non-magic filter. Blatent plug :-) I (B.A.McCauley@bham.ac.uk) have a generic magic filter bash script that selects the right filter to use based on the output of the `file` command. With my magic *dvi* filter (and existing separate *dvi* and *ps* filters) I can do things like:

```
lpr -d file1.dvi file2.div.Z file3.ps file4.ps.z
```

This is not yet on any anon-ftp sites so e-mail me for a copy. Once it's stable and I've had a bit of (positive) feedback I'll upload it.

Magic filters should never specified as `of` as the output filter only gets called once if a number of files are printed without a gap.

Magic filters as `if` can also be a problem as they may prevent you, say, *listing* a PostScript or nrff file.

17 Linux-FAQ: UUCP/News/Mail for Linux

This section is an expert of the UUCP-NEWS-MAIL-FAQ, posted twice monthly in c.o.l.a. This section and the complete FAQ about UUCP and others is written by Vince Skahan (vince@victoria.sea.wa.us), mail him for any corrections and information.

Last update June 1993.

17.1 General Information

17.1.1 What uucp, news, and mail software works under Linux?

Just about everything, most with no (or minimal) patching required...

The following is a quick but not all-inclusive list of software known to work under Linux:

UUCP

- Taylor UUCP 1.04

News transport

- Cnews 12/22/91 patchdate (no *optional* speedups)
- Cnews *performance release* (needs patching)
- INN 1.4 (needs patching) dynafeed
- nntp 1.5.11 *reference* release, slurp 1.05

News readers

- tin-1.2-p0
- trn-2.5 (probably also 3.0)
- nn-6.4.18
- tass
- rn4.4
- waffle 1.65

Mail transport

- smail2.5
- smail3.1.28
- sendmail5.65c+IDA

Mail program

- elm2.4.21 (probably also 2.4.22)
- pine
- procmail
- deliver

17.1.2 So where are the sources?

At your friendly neighborhood archive site where you find sources normally. Consultarchie for a list.

Config files for the portions of UUCP/News/Mail currently in the SLS distribution of Linux can be found on `sunsite.unc.edu` in the file `/pub/Linux/system/Mail/newspak-1.4.tar`.

17.1.3 Do I really need the sources?

Yes, if you want to *really* understand how the pieces are configured and how they fit together. If you already know how... probably not unless you want to change the paths used or do custom patching for some reason.

There is considerable documentation in the sources that doesn't get *installed* that helps fill in the blanks between what you need to know, and what you can figure out.

17.1.4 How do I get my system a news/mail/uucp feed?

Please see the `how to become a USENET site` periodic posting in `news.answers`.

17.2 UUCP

17.2.1 Why do I get 'timeout' on connections when I upgraded to uucp-1.04?

If you use a *Direct* device in the Devices file, there's now a 10 second timeout compiled in. Make the name of the Device anything other than *Direct*.

17.2.2 Why doesn't anonymous uucp seem to work?

The SLS anonymous uucp only works in Taylor mode because it's compiled with `HAVE_TAYLOR_CONFIG`. If you want to do anon uucp in `HDB` mode, you'll have to recompile the sources with just `HDB` defined.

`HDB` mode also seems rather particular about things like blank spaces and blank lines in your Permissions file. If you get log file entries that indicate 'no permission' for anon uucp, a good place to start is to make sure that there are no blank lines (comment 'em out).

17.3 News

17.3.1 What newsreader is best?

No way I'll open that can of worms other than to say that you definitely want something with "kill files" and probably want something with *threads*.

For what little it's worth, I use `nn`.

17.3.2 with `trn` — why do I get the message `dbinit` failed?

Because you didn't run `mthreads` from cron to keep the threads database up-to-date.

17.3.3 Why is my `nn` database getting continuously corrupted?

If you built it from sources and you asked for the `words` method of building the database, it will definitely get corrupted if your filesystem doesn't handle long filenames (ie...`comp.os.linux.announce` is more than 14-characters). Build it with the `by-number` way of creating the database and you'll be fine.

17.3.4 Why can't I post to moderated groups?

Probably because the newsreader is trying to call `/bin/mail` to send the mail and it doesn't like it. Replace the `/bin/mail` in SLS with the port of `mailx-5.3a.tar.z` from the `/pub/Linux/system/Mail` dir on `sunsite.unc.edu` and make it mode 2755 and group mail (like `elm`) and you'll be all set.

Another possibility is that you have a moderated newsgroup set up on your local system as not-moderated and somebody upstream is quietly deleting the article (some system's software, not a person). Make sure you run a `checkgroups` every now and then when the `checkgroups` article rolls by in `news.admin` every few weeks.

17.3.5 Why do I have problems that appear to be permission-related?

All the stuff in `/usr/local/lib/news` should be `news.news` except `/usr/local/lib/news/setnewsids` which should be setuid root.

Past versions of SLS had no `news` group or user by default, which caused lots of problems for people if they installed the news software *before* they added the `news` group and user.

17.3.6 Why doesn't SLS have the *Performance Release of C-news?*] INN? NNTP?

Because SLS can't have *everything*. It has the pre-performance C-news because it's a very fast and stable implementation for non-nntp sites that fits the needs of a normal uucp-only site. If you're that into having the most current of everything, you should probably be building it from the sources.

17.3.7 Why does doexpire say severe space shortage, when there's lots of room?

Because it can't read `/etc/mtab`. Make it mode 644. This happens when you unmount a mounted filesystem by root with a umask that doesn't permit world-read of files owned by root.

17.3.8 How do I set up NNTP?

Please see the extended description from Matt Welsh in the UUCP-NEWS-MAIL-FAQ that appears routinely in `c.o.l.a.`

17.4 Mail

17.4.1 Why is my `hostname.domainname` not in my headers and something funky's there?

Because you didn't edit one or more of the following files like the `README.linux` files in the `lib` directories told you to:

```
/usr/local/lib/smail/paths
/usr/local/lib/smail/config
/usr/local/lib/elm/elm.rc
```

17.4.2 Why can't elm write to my mailbox?

You need `elm` and `filter` to be mode 2755 and group `mail` and also have `/usr/spool/mail` mode 775 and group `mail`.

17.4.3 Is sendmail+IDA available for Linux instead of smail?

Yes !!! I've recently switched from `smail` to `sendmail+IDA` from the binary distribution on `sunsite.unc.edu` and it works great...and I'm a uucp-only site. Look in `/pub/Linux/system/Mail` for it...