

Linux Frequently Asked Questions with Answers (Part 3 of 6)

Source: <http://linux.derkeiler.com/Newsgroups/comp.os.linux.misc/2005-07/0407.html>

rkiesling_at_mainmatter.com

Date: 07/06/05

Date: 06 Jul 2005 04:31:02 GMT

Archive-Name: linux/faq/part3

URL: <http://www.mainmatter.com/>

Reply-to: rkiesling@mainmatter.com

Posting-Frequency: weekly

Last-modified: 12/04/2001

It's recommend that you create at least one Linux partition entirely under the 1024 logical cylinder limit, and boot from that. The other partitions will then be okay.

Also there seems to be a bit of trouble with the newer Ultra-DMA drives. I haven't gotten the straight scoop on them—but they are becoming a very common problem at the SVLUG installfests. When you can get 8 to 12 Gig drives for \$200 to \$300 it's no wonder.

[Jim Dennis]

4.2. How To Undelete Files.

In general, this is very hard to do on unices because of their multitasking nature. Undelete functionality for the ext2fs file system is being worked on, but don't hold your breath.

There are a number of packages available which instead provide new commands for deleting and copying which move deleted files into a 'wastebasket' directory. The files can be recovered until cleaned out automatically by background processing.

The Midnight Commander file manager provides an undelete facility that uses Ext2 file system library functions and an undelete directory for each file system. Commercial distribution packages of MC may or may not have this feature enabled, so be sure to look in the source code distribution for instructions on how to enable the undelete feature.

Alternatively, you can search the raw disk device which holds the file system in question. This is hard work, and you will need to be logged

in as root to do this. But it can be done. Run grep on the raw device;
e.g.:

```
grep -b 'bookmarks' /dev/hda
```

If the data has not been overwritten, you should be able to recover it with a text editor.

[Dave Cinege, Daniel Novotny]

4.3. How To Make Backups.

You can back up a directory hierarchy or complete file system to any media using GNU tar or cpio, the standard *nix tools for this purpose. tar seems to be the more commonly used program currently, and includes command line options to make compressed, incremental, and multi-volume backups. Complete information is contained in the documentation, which is in GNU Texinfo format.

The free program, Amanda, receives a lot of mentions on Usenet. Its home page is <http://www.amanda.org/>.

Several commercial backup utilities also exist. They are often included in commercial distributions.

4.4. How To Resize a Partition (Non-Destructively).

Use the FIPS.EXE program, included with most Linux distributions, under MS-DOS.

GNU parted, a partition editor, is stable enough for non-guru, mere-mortal use with relative confidence. Source code for the latest version is at: <ftp://ftp.gnu.org/pub/gnu/parted/>. There's also a boot disk image for resizing root partitions and for running parted on non-Linux machines. The disk image may be easier for beginners. Building from source could require some extra configuration.

Parted also has tutorial-style, plain-text documentation for Linux and FAT (MS-DOS) file systems.

Also, some commercial distributions come with their own partitioning software, like Partition Magic.

4.5. Is There a Defragmenter for Ext2fs?

Yes. There is defrag, a Linux file system defragmenter for ext2, Minix, and old-style ext file systems. It is available at <ftp://metalab.unc.edu/pub/Linux/system/filesystems/defrag-0.70.tar.gz>.

Users of the ext2 file system can probably do without defrag, because ext2 contains extra code to keep fragmentation reduced even in very

full file systems.

4.6. How To Create a File System on a Floppy.

To format a 3.5-inch, high density floppy:

```
$ fdformat /dev/fd0H1440
$ mkfs -t ext2 -m 0 /dev/fd0H1440 1440
```

For a 5.25 inch floppy, use fd0h1200 and 1200 as appropriate. For the B: drive use fd1 instead of fd0.

The `-m 0` option tells `mkfs.ext2` not to reserve any space on the disk for the superuser—usually the last 10% is reserved for root.

The first command performs a low-level format. The second creates an empty file system. You can mount the floppy like a hard disk partition and simply `cp` and `mv` files, etc.

Device naming conventions generally are the same as for other unices. They can be found in Matt Welsh's Installation and Getting Started guide. Refer to ("Where Is the Documentation?") A more detailed and technical description is Linux Allocated Devices by H. Peter Anvin, hpa@zytor.com, which is included in LaTeX and ASCII form in the kernel source distribution (probably in `/usr/src/kernel/Documentation/`), as `devices.tex` and `devices.txt`.

4.7. Does Linux Support Virtualized File Systems Like RAID?

The most recent Linux kernels support software RAID, and they will work with RAID disk controllers.

An automounter for NFS partitions is part of most Linux distributions.

In addition, several virtual file system projects exist. One of them, the Linux Logical Volume Manager, is located at <http://linux.msede.com/lvm/>.

4.8. Does Linux Support File System Encryption?

Yes. One file system, `ppdd`, is archived at <http://pweb.de.uu.net/flexsys.mtk/>.

4.9. Linux Prints Nasty Messages about Inodes, Blocks, and the Like.

You may have a corrupted file system, probably caused by not shutting Linux down properly before turning off the power or resetting. You need to use a recent shutdown program to do this—for example, the one included in the `util-linux` package, available on `sunsite` and `tsx-11`.

If you're lucky, the program fsck (or e2fsck or xfsck as appropriate if you don't have the automatic fsck front-end) will be able to repair your file system. If you're unlucky, the file system is trashed, and you'll have to re-initialize it with mkfs (or mke2fs, mkxfs, etc.), and restore from a backup.

NB: don't try to check a file system that's mounted read/write—this includes the root partition, if you don't see

```
VFS: mounted root ... read-only
```

at boot time.

4.10. The Swap Area Isn't Working.

When you boot (or enable swapping manually) you should see

```
Adding Swap: NNNNk swap-space
```

If you don't see any messages at all you are probably missing

```
swapon -av
```

(the command to enable swapping) in your /etc/rc.local or /etc/rc.d/* (the system startup scripts), or have forgotten to make the right entry in /etc/fstab:

```
/dev/hda2 none swap sw
```

for example.

If you see:

```
Unable to find swap-space signature
```

you have forgotten to run mkswap. See the manual page for details; it works much like mkfs.

Running, free in addition to showing free memory, should display:

```
total used free  
Swap: 10188 2960 7228
```

If typing "cat /proc/swaps" reveals only file or partition names, but no swap space information, then the swap file or partition needs re-initialization.

Use fdisk (as root) to determine which partition on a hard drive has been designated as the swap partition. The partition still needs to be initialized with mkswap before enabling it with swapon.

[Andy Jefferson, Steve Withers]

4.11. How To Add Temporary Swap Space.

In addition to a swap partition, Linux can also use a swap file. Some programs, like g++, can use huge amounts of virtual memory, requiring the temporary creation of extra space. To install an extra 64 MB of swap space, for example, use the following shell commands:

```
# dd if=/dev/zero of=/swap bs=1024 count=65535
# mkswap /swap
# swapon /swap
```

The count= argument to dd determines how big the swap file will be. In this example the swap file's name is /swap, but the name and location are, generally, arbitrary, depending only on the file system's available space and your having write permissions in the directory.

When you don't need the swap space any more, remove it with the following statements:

```
# swapoff /swap
# rm /swap
```

Take a look also at the Installation HOWTO and Installation & Getting Started for detailed instructions.

If that still doesn't provide enough swap space, refer to ("How To Have More Than 128Mb of Swap.")

4.12. How To Remove LILO So the System Boots DOS Again?

The lilo program (not the complete LILO package), uses the command line option -u to uninstall the LILO boot loader. You have to supply the device name of the device you installed LILO on, for example:

```
lilo -u /dev/hda
```

This rewrites the original, pre-LILO master boot record back to the first hard drive, from the boot record saved in /boot/boot.0300. If you installed LILO to a partition as a secondary boot loader, for example, /dev/hda1, lilo re-installs the original boot sector from the save file /boot/boot.0301. Refer to the lilo manual page for details. Thanks to Villy Kruse for reminding me to update this answer.

If you have an earlier version of LILO, you will have to use the DOS (MS-DOS 5.0 or later, or OS/2) FDISK /MBR (which is not documented). This will restore a standard MS-DOS Master Boot Record. If you have DR-DOS 6.0, go into FDISK.EXE in the normal way and then select the Re-write Master Boot Record option.

If you create a boot floppy during the Windows installation process, make sure that it contains the programs FDISK.EXE, FORMAT.COM, and SYS.COM, and use that to re-install MS-DOS on the hard disk.

If you don't have MS-DOS or DR-DOS, you need to have the boot sector that LILO saved when you first installed it. You did keep that file, didn't you? It's probably called boot.0301 or some such. Type:

```
dd if=boot.0301 of=/dev/hda bs=445 count=1
```

(or /dev/sda if you're using a SCSI disk). This may also wipe out your partition table, so beware! If you're desperate, you could use

```
dd if=/dev/zero of=/dev/hda bs=512 count=1
```

This will erase your partition table and boot sector completely: you can then reformat the disk using your favorite software. But this will render the contents of your disk inaccessible—you'll lose it all unless you're an expert.

Note that the DOS MBR boots whichever (single!) partition is flagged as "active." You may need to use fdisk to set and clear the active flags on partitions appropriately.

4.13. Why Does fdformat Require Superuser Privileges?

The system call to format a floppy can only be done as root, regardless of the permissions of /dev/fd0*. If you want any user to be able to format a floppy, try getting the fdformat2 program. This works around the problems by being setuid to root.

4.14. The System Checks the Ext2fs Partitions Each Reboot.

Refer to ("EXT2-fs: warning: mounting unchecked file system.")

4.15. Root File System Is Read-Only.

Remount it. If /etc/fstab is correct, you can simply type:

```
mount -n -o remount /
```

If /etc/fstab is wrong, you must give the device name and possibly the type, too: e.g.

```
mount -n -o remount -t ext2 /dev/hda2 /
```

To understand how you got into this state, see, ("EXT2-fs: warning: mounting unchecked file system.")

4.16. What Is /proc/kcore?

None of the files in /proc are really there—they're all, "pretend," files made up by the kernel, to give you information about the system and don't take up any hard disk space.

/proc/kcore is like an "alias" for the memory in your computer. Its size is the same as the amount of RAM you have, and if you read it as a file, the kernel does memory reads.

4.17. The AHA1542C Doesn't Work with Linux.

The option to allow disks with more than 1024 cylinders, which the AHA1542C card can recognize, is only required as a workaround for a PC-compatible BIOS misfeature and should be turned `off` under Linux. For older Linux kernels you need to turn off most of the `advanced BIOS` options—all but the one about scanning the bus for bootable devices.

4.18. Where Is the Journalling File System on the Net?

The journalling file system, named Reiserfs has just been released from testing. It is said to make Linux even faster than Linux with the Ext2 file system installed. Complete information is available at <http://devlinux.org/namesys/>.

5. Porting, Compiling and Obtaining Programs

5.1. How To Compile Programs.

Most Linux software is written in C and compiled with the GNU C compiler. GCC is a part of every Linux distribution. The latest compiler version, documentation, and patches are on <ftp://ftp.gnu.org/pub/gnu/>.

Programs that are written in C++ must be compiled with the GNU G++ compiler, which is also included in Linux distributions and available from the same place as GCC.

To build version 2.0.x kernels, you will need GCC version 2.7.2.x, approximately. Trying to build an early Linux kernel with a different compiler, like GCC 2.8.x, EGCS, or PGCC, may cause problems because of GCC related code dependencies. Kernel versions 2.3 and 2.4 should compile correctly with more recent compilers.

Information on the EGCS compiler is at <http://www.gnu.org/software/gcc/gcc.html>.

Note that at this time, the kernel developers are not answering bug requests for earlier kernels, but instead are concentrating on developing 2.4.x version kernels and maintaining 2.2.x version kernels.

[J.H.M. Dassen, Axel Boldt]

5.2. How To Install GNU Software.

On a correctly configured system, installing a GNU software package requires four steps:

- * With the source.tar.gz archive in the /usr/src/ directory, or wherever you maintain your source files, untar and decompress the package with the command:

```
tar zxvf package-name.tar.gz
```

- * Run the ./configure script in the untarred source archive's top-level directory with whatever command line arguments you need. The options that configure recognizes are usually contained in a file called INSTALL or README.
- * Run make. This will build the source code into an executable program (or programs) and may take a few minutes or a few hours, depending on the speed of the computer and the size of the package.
- * Run make install. This will install the compiled binaries, configuration files, and any libraries in the appropriate directories.

5.3. Where To Get Java.

The Sun Microsystems Java runtime environments and developer's kits are at <http://java.sun.com/java/>.

You can also get the source code, which is licensed by Sun Microsystems.

The home page of IBM's Jikes Java compiler is <http://www10.software.ibm.com/developerworks/opensource/jikes>.

There is a version of Sun's HotJava browser for Linux at: <http://www.java.sun.com/products/hotjava/>.

Kaffee, a free Java clone, is available from: <http://www.kaffe.org/>.

There is a resource page of free and commercial Java packages at: <http://www.blackdown.org/java-linux.html>.

Netscape Communicator comes with its own version of the Java Runtime Environment, so if you need Java only to view Web graphics, you may already have the runtime version of the Java Virtual Machine and libraries that you need installed on your system. You can download Communicator from <http://www.netscape.com/>.

5.4. How To Port XXX to Linux.

In general, *nix programs need very little porting. Simply follow the installation instructions. If you don't know—and don't know how to find out—the answers to some of the questions asked during the installation procedure, you can guess, but this tends to produce buggy programs. In this case, you're probably better off asking someone else to do the port. If you have a BSD-ish program, you should try using `-I/usr/include/bsd` and `-lbsd` on the appropriate parts of the compilation lines.

5.5. What Is ld.so and How To Get It?

ld.so is the dynamic library loader. Each binary using shared libraries used to have about 3K of start-up code to find and load the shared libraries. Now that code has been put in a special shared library, `/lib/ld.so`, where all binaries can look for it, so that it wastes less disk space, and can be upgraded more easily. ld.so can be obtained from <http://tsx-11.mit.edu/pub/linux/packages/GCC/> and mirror sites. The latest version at the time of writing is `ld.so.1.9.5.tar.gz`. `/lib/ld-linux.so.1` is the same thing for ELF ("What's all this about ELF? ") and comes in the same package as the a.out loader.

5.6. How To Upgrade the Libraries without Trashing the System.

Warning

Note: You should always have a rescue disk set ready when you perform this procedure, in the likely event that something goes wrong!

This procedure is especially difficult if you're upgrading very old libraries like libc4. But you should be able to keep libc4 on the same system with libc5 libraries for the programs that still need them. The same holds true for upgrading from libc5 to the newer-yet glibc2 libraries.

The problem with upgrading dynamic libraries is that, the moment you remove the old libraries, the utilities that you need to upgrade to the new version of the libraries don't work. There are ways around around this. One is to temporarily place a spare copy of the run time libraries, which are in `/lib/`, in `/usr/lib/`, or `/usr/local/lib/`, or another directory that is listed in the `/etc/ld.so.conf` file.

For example, when upgrading libc5 libraries, the files in `/lib/` might look something like:

```
libc.so.5
libc.so.5.4.33
libm.so.5
libm.so.5.0.9
```

These are the C libraries and the math libraries. Copy them to another directory that is listed in `/etc/ld.so.conf`, like `/usr/lib/`:

```
$ cp -df /lib/libc.so.5* /usr/lib/  
$ cp -df /lib/libm.so.5* /usr/lib/  
$ ldconfig
```

Be sure to run `ldconfig` to upgrade the library configuration.

The files `libc.so.5` and `libm.so.5` are symbolic links to the actual library files. When you upgrade, the new links will not be created if the old links are still there, unless you use the `-f` flag with `cp`. The `-d` flag to `cp` will copy the symbolic link itself, and not the file it points to.

If you need to overwrite the link to the library directly, use the `-f` flag with `ln`.

For example, to copy new libraries over the old ones, try this. Make a symbolic link to the new libraries first, then copy both the libraries and the links to `/lib/`, with the following commands.

```
$ ln -sf ./libm.so.5.0.48 libm.so.5  
$ ln -sf ./libc.so.5.0.48 libc.so.5  
$ cp -df libm.so.5* /lib  
$ cp -df libc.so.5* /lib
```

Again, remember to run `ldconfig` after you copy the libraries.

If you are satisfied that everything is working correctly, you can remove the temporary copies of the old libraries from `/usr/lib/` or wherever you copied them.

5.7. How To Use Code or a Compiler Compiled for a 486 on a 386.

Yes, unless it's the kernel.

The `-m486` option to GCC, which is used to compile binaries for x486 machines, merely changes certain optimizations. This makes for slightly larger binaries that run somewhat faster on a 486. They still work fine on a 386, though, with a small performance hit.

However, from version 1.3.35 the kernel uses 486 or Pentium-specific instructions if configured for a 486 or Pentium, thus making it unusable on a 386.

GCC can be configured for a 386 or 486; the only difference is that configuring it for a 386 makes `-m386` the default and configuring for a 486 makes `-m486` the default. In either case, these can be overridden on a per-compilation basis or by editing `/usr/lib/gcc-lib/i*-linux/n.n.n/specs`.

There is an alpha version of GCC that knows how to do optimization well for the 586, but it is quite unreliable, especially at high optimization settings. The Pentium GCC can be found on <ftp://tsx-11.mit.edu/pub/linux/ALPHA/pentium-gcc/>. The ordinary 486 GCC supposedly produces better code for the Pentium using the `-m386`, or at least slightly smaller.

5.8. What Does "gcc -O6" Do?

Currently, the same as `-O2` (GCC 2.5) or `-O3` (GCC 2.6, 2.7). Any number greater than that does the same thing. The Makefiles of newer kernels use `-O2`, and you should probably do the same.

5.9. Where Are linux/*.h and asm/*.h?

The files `/usr/include/linux/` and `/usr/include/asm/` are often soft links to the directories where the kernel headers are. They are usually under `/usr/src/kernel*/`.

If you don't have the kernel sources, download them. Refer to the answer: ("How To Upgrade/Recompile a Kernel.")

Then, use `rm` to remove any garbage, and `ln` to create the links:

```
$ rm -rf /usr/include/linux /usr/include/asm
$ ln -sf /usr/src/linux/include/linux /usr/include/linux
$ ln -sf /usr/src/linux/include/asm-<architecture> /usr/include/asm
```

The assembly language files reside in architecture-specific directories, so you need to link `/usr/src/include/asm` to `/usr/src/linux/include/asm-i386` on PC compatible systems, to `/usr/src/linux/include/asm-sparc` on Sun Sparc systems, to `/usr/src/linux/include/asm-ppc` on PPC systems, and so on.

You'll also find that you may need to do ``make config'` as in a newly-unpacked kernel source tree, to create `linux/autoconf.h`.

5.10. What To Do about Errors Trying to Compile the Kernel.

See the previous question regarding the header files.

Remember that when you apply a patch to the kernel, you must use the `"-p0"` or `"-p1"` option: otherwise, the patch may be misapplied. See the patch manual page for details.

"ld: unrecognized option ``-qmagic'`" means that you should get a newer linker, from <ftp://tsx-11.mit.edu/pub/linux/packages/GCC/>, in the file `binutils-2.8.1.0.1.bin.tar.gz`.

5.11. How To Make a Shared Library.

For ELF,

```
$ gcc -fPIC -c *.c  
$ gcc -shared -Wl,-soname,libfoo.so.1 -o libfoo.so.1.0 *.o
```

For a.out, get `tools-n.nn.tar.gz` from tsx-11.mit.edu/pub/linux/packages/GCC/src/. It comes with documentation that will tell you what to do. Note that a.out shared libraries are a very tricky business. Consider upgrading your libraries to ELF shared libraries. See the ELF HOWTO, at [ftp://metalab.unc.edu/pub/Linux/docs/HOWTO/](http://metalab.unc.edu/pub/Linux/docs/HOWTO/)

5.12. Programs Are Very Large.

With an ELF compiler ("What's All This about ELF? glibc?"), the most common cause of large executables is the lack of an appropriate `.so` library link for one of the libraries you're using. There should be a link like `libc.so` for every library like `libc.so.5.2.18`.

With an a.out compiler the most common cause of large executables is the `-g` linker (compiler) flag. This produces (as well as debugging information in the output file) a program which is statically linked—one which includes a copy of the C library instead of a dynamically linked copy.

Other things worth investigating are `-O` and `-O2`, which enable optimization (check the GCC documentation), and `-s` (or the `strip` command) which strip the symbol information from the resulting binary (making debugging totally impossible).

You may wish to use `-N` on very small executables (less than 8K with the `-N`), but you shouldn't do this unless you understand its performance implications, and definitely never with daemons.

5.13. Does Linux Support Threads or Lightweight Processes?

As well as the Unix multiprocessing model involving heavyweight processes, which is of course part of the standard Linux kernel, there are several implementations of lightweight processes or threads. Recent kernels implement a thread model, `kthreads`. In addition, there are the following packages available for Linux.

- * GNU `glibc2` for Linux has optional support for threads. The archive is available from the same place as `glibc2`, [ftp://ftp.gnu.org/pub/gnu/](http://ftp.gnu.org/pub/gnu/)

- * In `sipb.mit.edu:/pub/pthread/` or `ftp.ibp.fr:/pub/unix/threads/pthreads`. Documentation isn't in the package, but is available on the World Wide Web at http://www.mit.edu:8001/people/proven/home_page.html. Newer Linux `libc`'s contain the `pthreads` source. The GNU Ada compiler on [ftp://metalab.unc.edu/pub/Linux/devel/lang/ada/gnat-3.01-linux+elf](http://metalab.unc.edu/pub/Linux/devel/lang/ada/gnat-3.01-linux+elf)

.tar.gz contains binaries made from that source code.

- * In <ftp://ftp.cs.washington.edu/pub/qt-001.tar.Z> is QuickThreads. More information can be found in the technical report, available on the same site is /tr/1993/05/UW-CSE-93-05-06.PS.Z.
- * In <http://gummo.doc.ic.ac.uk/rex/> is lwp, a very minimal implementation.
- * In <ftp://ftp.cs.fsu.edu/pub/PART/>, an Ada implementation. This is useful mainly because it has a lot of Postscript papers that you'll find useful in learning more about threads. This is not directly usable under Linux.

Please contact the authors of the packages in question for details.

5.14. Where To Find lint for Linux.

Roughly equivalent functionality is built into GCC. Use the `-Wall` option to turn on most of the useful extra warnings. See the GCC manual for more details (type `F1-i` in Emacs and select the entry for GCC).

There is a freely available program called `lclint` that does much the same thing as traditional `lint`. The announcement and source code are available at on <ftp://larch.lcs.mit.edu/pub/Larch/lclint/>; on the World Wide Web, look at <http://lclint.cs.virginia.edu/>.

5.15. Where To Find Kermit for Linux.

Kermit is distributed under a non-GPL copyright that makes its terms of distribution somewhat different. The sources and some binaries are available on <ftp://kermit.columbia.edu>.

The WWW Home Page of the Columbia University Kermit project is <http://www.columbia.edu/kermit/>.

5.16. How To Use Linux with a Cable Modem.

The [www.CablemodemInfo.com](http://www.cablemodeminfo.com) and xDSL Web page at <http://www.cablemodeminfo.com/> has a section devoted to Linux.

5.17. Is There an ICQ Program That Runs under Linux?

Several ICQ clients are available on metalab.unc.edu. (Refer to: "Where Are the Linux FTP Archives?") ICQ itself does not have a Linux client, but there is a Java client at <http://www.mirabilis.com/download/>.

6. Solutions to Common Miscellaneous Problems

6.1. FTP Transfers Seem to Hang.

FTP transfers that die suddenly are due, apparently, to some form of overrunning buffer. It occurs both with Linux and Microsoft servers.

On Linux systems, The problem seems to occur most commonly with the distribution's server software.

If you receive ftp: connection refused errors, then the problem is likely due to a lack of authentication. Refer to "FTP or Telnet Server Won't Allow Logins.."

One remedy is to be replacing the distribution FTP server with the Linux port of the OpenBSD FTP server. The home page is:
<http://www.eleves.ens.fr:8080/home/madore/programs/>

To install the BSD server, follow the installation instructions, and refer to the manual pages for inetd and inetd.conf. (If you have the newer xinetd, see below.) Be sure to tell inetd to run the BSD daemon alone, not as a subprocess of, for example, tcpd. Comment out the line that begins "ftp" in the /etc/inetd.conf file and replace it with a line similar to (if you install the new ftpd in /usr/local/sbin/):

```
# Original entry, commented out.  
#ftp stream tcp nowait root /usr/sbin/tcpd /usr/sbin/in.ftpd
```

```
# Replacement entry:  
ftp stream tcp nowait root /usr/local/sbin/ftpd -l
```

The replacement daemon will become effective after rebooting or sending (as root) a SIGHUP to inetd, e.g.:

```
# kill -HUP inetd
```

To configure xinetd, create an entry in /etc/xinetd.d per the instructions in the xinetd.conf manual page. Make sure, again, that the command-line arguments for ftpd are correct, and that you have installed the /etc/ftpusers and /etc/pam.d/ftp files. Then restart xinetd with the command: /etc/rc.d/init.d/xinetd restart. The command should report "OK," and the restart will be noted in the system message log.

6.2. Free Dumps Core.

In Linux 1.3.57 and later, the format of /proc/meminfo was changed in a way that the implementation of free doesn't understand. Get the latest version, from metalab.unc.edu, in /pub/Linux/system/Status/ps/procps-0.99.tgz.

6.3. Netscape Crashes Frequently.

Netscape shouldn't crash, if it and the network are properly configured. Some things to check:

- * Make sure that the MOZILLA_HOME environment variable is correctly set. If you installed Netscape under /usr/local/netscape/, for

example, that should be the value of MOZILLA_HOME. Set it from the command line (e.g. "export MOZILLA_HOME="/usr/local/netscape"" under bash or add it to one your personal or system initialization files. Refer to the manual page for your shell for details.

- * If you have a brand-new version of Netscape, try a previous version, in case the run-time libraries are slightly incompatible. For example, if Netscape version 4.75 is installed (type "netscape --version" at the shell prompt), try installing version 4.7. All versions are archived at <ftp://ftp.netscape.com/>.
- * Netscape uses its own Motif and Java Runtime Environment libraries. If a separate version of either is installed on your system, ensure that they aren't interfering with Netscape's libraries; e.g., by un-installing them.
- * Make sure that Netscape can connect to its default name servers. The program will appear to freeze and time out after several minutes if it can't. This indicates a problem with the system's Internet connection; likely, the system can't connect to other sites, either.

6.4. FTP or Telnet Server Won't Allow Logins.

This applies to server daemons that respond to clients, but don't allow logins. On new systems that have Pluggable Authentication Modules installed, look for a file named, "ftp," or "telnet," in the directory /etc/pam/ or /etc/pam.d/. If the corresponding authentication file doesn't exist, the instructions for configuring FTP and Telnet authentication and other PAM configuration, should be in /usr/doc/pam-<version>. Refer also to the answer for "FTP server says: "421 service not available, remote server has closed connection."."

If it's an FTP server on an older system, make sure that the account exists in /etc/passwd, especially "anonymous."

This type of problem may also be caused a failure to resolve the host addresses properly, especially if using Reverse Address Resolution Protocol (RARP). The simple answer to this is to list all relevant host names and IP addresses in the /etc/hosts files on each machine. (Refer to the example /etc/hosts and /etc/resolv.conf files in: "Sendmail Pauses for Up to a Minute at Each Command..") If the network has an internal DNS, make sure that each host can resolve network addresses using it.

If the host machine doesn't respond to FTP or Telnet clients at all, then the server daemon is not installed correctly, or at all. Refer to the manual pages: inetd and inetd.conf on older systems, or xinetd and xinetd.conf, as well as ftpd, and telnetd.

6.5. How To Keep Track of Bookmarks in Netscape?

This probably applies to most other browsers, too. In the Preferences/Navigator menu, set your home page to Netscape's bookmarks.html file, which is located in the .netscape (with a leading period) subdirectory. For example, if your login name is "smith," set the home page to:

<file://home/smith/.netscape/bookmarks.html>

Setting up your personal home page like this will present you with a nicely formatted (albeit possibly long) page of bookmarks when Netscape starts. And the file is automatically updated whenever you add, delete, or visit a bookmarked site.