

PPC-E5 Flash Linux

This section refers to the EMAC distribution of Linux, which is currently available for the PPC-E5 on a 16MB DOC Flash Disk.

Introduction

The EMAC Panel PC (PPC) is designed to act as a general-purpose embedded Linux server. The PPC runs the Linux operating system kernel (a UNIX work-alike), and as such retains all of the networking and communications capabilities one would expect of such an operating system. The PPC v2.0 includes Linux kernel 2.2.14 and GNU standard C libraries version 2.1.3 (glibc6 v2.2.13).

Right out of the box, the PPC is configured for ethernet, serial IP, serial terminal, and raw serial connections. It includes a web server, an FTP server, a TELNET server, and a periodic command scheduler. Other devices, including PPP modem links, may be added and configured by the user.

This section is designed to help you configure and work with the Linux system, and understand what you do and why. The first section of this section, explains what the PPC can do, how it does it, and how each family of utilities can be set up and configured. The second section of this manual, Configuration, explains the details of setting up the PPC using the EMAC configuration utility. It is hoped that this format will first help the user understand what to configure and why, before explaining exactly how. Users with thorough knowledge of Linux should at least skim section one briefly before moving on to section two, as section one contains important information about differences between the EMAC PPC and other (full-sized) Linux distributions.

The Boot Process

EMAC's flash linux uses a specially patched version of the LILO bootloader to boot an X windows terminal on power up. When the system is first booted, the flash disk is mounted read-only for corruption resistance. Mounting the flash read-write is a two-step process: First, open an Aterm window, issue the *su* command and enter the root password to gain superuser privileges. Next, issue the command *mount -o remount,rw /dev/diskroot /*. The default root password for the PPC-E5 is *emac_inc*.

The DOC partition

The DOC is a device file of type fla. This is a non-standard device driver that is included with the EMAC distribution specifically for mounting the DOC. The entire filesystem is mounted on the first partition, */dev/fla1*. Temporary space is provided in a ramdisk, which is mounted as */var*. No swap space is included, as excessive writes will cause the DiskOnChip device to wear.

The Start Window

Probably the most important basic feature to know is the start window. The start window pops up any time an area outside of an open window is touched. It presents the user with 4 options:

1. *Desktop* - desktop configuration options; change the look of your desktop, update the start menu etc.
2. *Windows* - windows options for resizing, moving, closing etc.
3. *Aterm* - A bash command prompt (keyboard required and monitor recommended)
4. *Contrast/Volume* - Scrollbars for setting the contrast and volume

Aterm

Aterm is a virtual terminal that is available as an option on the start menu. It is a smaller equivalent of the common Xterm. Once it is selected a window is opened which presents the user with a bash prompt. To save space, all common functions are included in two multi-call binaries called *Busybox* and *Tinylogin*. To see a list of available commands, type *busybox* or *tinylogin*. In order to invoke a command contained in BusyBox or Tinylogin, simply invoke the command you wish to use. Aterm allows the user easy access to the command line. The user can also shutdown X by simultaneously typing <ctrl><alt><backspace> with a keyboard connected to the PPC.

Touchscreen

The touchscreen is configured in `/usr/X11R6/lib/X11`. This config file contains all the options for the touchscreen. If the pointer does not follow your finger, this file probably needs adjusting. Check if these values are set correctly:

```
Min x = 1980
Max x = 40
Min y = 1960
Max y = 35
Swapxy
```

The first 4 locate what is to be displayed on the touchscreen with respect to the VGA display; the values here should set it appropriately. Swapxy rotates the touchscreens display by 90 degrees.

Files can be manipulated with the text editor `ae` that is included in the flash Linux distribution.

Contrast/Volume

A small program is included for changing the contrast and volume settings of the PPC-E5. To use it, first touch outside of any windows currently open to bring up the start menu. Select Contrast/Volume to bring up the adjustment scrollbars. The scrollbars select contrast and volume settings, which will be saved in non-volatile memory and restored whenever the system is re-booted. Once you are satisfied with the settings, the program can be closed like any other window by touching the X.

Part I: PPC Basics

Section 1: Users and Security

Linux follows the UNIX security model of separate users and username / password authentication. This means that each person to use the machine has a separate "account," distinguished by a username (also called a "login name"). To gain access to an account on a UNIX or Linux machine, a person must know both a username and the password associated with it.

The UNIX security model also includes the concept of groups. Each account (user) is a member of one or more groups, and each group has one or more members. Groups allow multiple people to access the same file(s), yet still prevent non-privileged users from viewing or modifying those files.

Each user (and each group) can have different permissions and privileges. These privileges take the form of "file permissions." Each file in a UNIX system has two ownerships. Every file is owned by a single user, and by a group. Each file also has three sets of three permissions: read, write, and execute permissions may be specified separately on each file for the file's user owner, the file's group owner, and everyone on the entire system. These permissions are expressed by the following bitfield:

User	Group	Everyone
read-write-execute	read-write-execute	read-write-execute

When describing permissions, this bitfield is expressed as an octal (base 8) number. For example, permissions 754 mean that the user owner of the file may read, write, or execute the file (7). The group may read or execute, but not write (5). Other users on the system may read the file, but not write it or execute it (4).

Nearly every important file on a UNIX system is owned by a single user and group, **root** (**root** user, **root** group). The **root** account is essential to the security of any UNIX system. Having all of the important executable and configurations files on the system owned by root means that only the root user may alter the way the system works. Other users have no permission to execute or change system files, and therefore they cannot modify or damage the system. As such, it is critical that the root password be kept secret, and changed as often as is practical. A clumsy or malicious individual with the PPC's **root** password could render the entire system inoperative in a matter of seconds. Every PPC ships from the factory with the same **root** password - this means that anyone who has ever purchased an PPC knows your **root** password! EMAC recommends you change the PPC's **root** password as soon as you've finished reading the manual.

Aside from the root account, which is present on every system, the PPC has one additional account configured. The username is **www**. This user does not have privileges to change the system configuration. Instead, this account is designed to allow a semi-privileged user to upload HTML files to the PPC. From there they may be served by the PPC's webserver. This account may also be used when transferring new executables to the PPC. While it is much more difficult for an individual with the **www** password to damage the PPC, it is still possible. EMAC therefore recommends you change **www**'s password right after you change the **root** password.

Section 2: Command-line Tools

Certain command-line utilities are present on nearly every UNIX system. The PPC provides as many of these as is practical in the space available. However, it is important to note that some of these condensed command line tools function somewhat differently than the versions you might find in a full-sized Linux distribution. If you are already familiar with tools such as `grep` and `sed`, you may find that the syntax differs slightly from what you are used to, or that certain switches are not supported.

The PPC's main shell is GNU `bash`, with full scripting capability. A lesser shell, `ash`, is also available. Most of the standard command-line tools are provided by `busybox`, a multi-call binary which does the work of many different programs, including `grep`, `sed`, `awk`, and `ash`. In some cases, we felt that the `busybox` implementation of a particular function was inadequate (`tar`, for example). In these cases, regular versions of these programs from full-sized distributions are used. Descriptions of some of the most common command line utilities follow. Keep in mind when running these programs that Linux (and UNIX in general) is case-sensitive; capitalization matters. In all cases, the program name must be types in all lowercase letters.

The `ls` utility works much like the `dir` command in DOS. With no parameters, `ls` will list the files in the current directory. A directory name may also be specified on the command line, `ls` will then list this directory instead. For more detail, the `-l` switch (no quotes) may be specified between `ls` and the directory name. This will display information about files including ownership, permissions, size, and modification times.

The `cd` command changes the current directory, exactly as in DOS. As UNIX has no concept of "drives," all files are kept in a single unified file system. Forward slashes (`/`) are used to separate subdirectory names.

Each user has a "home directory." This directory is where a user will begin after logging in. For the root user, this directory is `/root/`. For other (regular) users, this directory will usually be `/home/<username>/`. When using the command line, a user's home directory may be abbreviated by a tilde character (`~`). For example, `cd ~` will bring you to your home directory. The command `cd ~foo` will take you to user `foo`'s home directory (assuming you have permission to read that directory). Similarly, `~/foo/bar` will execute a program called `bar` in `foo`'s home directory. Note that a tilde abbreviation should never be prefixed with a `/`.

The `chown`, `chgrp`, and `chmod` commands modify file ownerships and permissions. The `chown` utility will change the owner of a file to one specified on the command line (`chown <user> <file>`). `Chgrp` will change the group owner of a file in the same manner. The `chmod` utility changes the permissions of a file (`chmod <permissions> <file>`). The new permissions are specified numerically, as described previously. Note that you must have write permission on a file to use these utilities on it.

The `cat` utility will display the entire contents of a file to the screen. This should only be done to text files, as "catting" binary files can cause the shell to enter odd states in which it prints unreadable characters. There is no way to pause the output from `cat`. The `more` utility fixes this shortcoming by displaying a specified file one page at a time. Simply press a key to view the next page of text. The `more` utility is the standard way to view text files on the PPC.

The PPC includes a simple text editor, called `ae`. It is invoked from the command line as `as <filename>`. When `ae` starts, its help screen is displayed on top of the edit window. If for some reason the help screen does not appear, hold down the control key and press `x`, then `h`. `Ae` is a "modeless" editor, like `notepad` or GNU `emacs` (in contrast to `vi`, which is a mode-based editor). As such, all functions will work all of the time, including the arrow keys and backspace.

To move a file in the PPC filesystem, use the mv command. Its syntax is identical to the DOS move command. Keep in mind that when the destination parameter ends in a /, that parameter will be treated as a directory, and the file will be moved to that directory (if it exists) and the name will remain unchanged. If the destination parameter does not end in a /, it will be treated as a file name, and the file's name will be changed. The cp command is likewise similar to the DOS copy utility. Linux does not have a rename function, but mv can be used to change a file's name. The rm command will delete a file. To delete a directory, the "-r" switch must be specified before the directory name to delete. To prevent rm from prompting for confirmation, use the "-f" switch. CAUTION: There is no undelete utility. File deleted from the PPC are gone, permanently.

The PPC includes a simple terminal program, minicom. Minicom works much like procomm or hyperterminal. Minicom may be used to communicate with serial devices, including embedded devices and terminal servers. It can be used to talk to modems and debug new hardware. Minicom can use a configuration file, but does not have one from the factory -- if no configuration file is present, minicom will start with its defaults. Once minicom is started, press Control-A, Z to display a list of options. From this menu you may configure the baud rate, the terminal type settings, scripting, and many other parameters. You can also use the "o" option of this menu to configure and create a configuration script. Minicom will use the device /dev/modem by default, which is linked to /dev/ttyS2. This may also be changed using the "o" option. A detailed description of the configuration and use of minicom is beyond the scope of this document.

Section 3: Loadable Modules and Device Drivers

There are two ways to use a device driver in Linux: it can be compiled into the kernel, or it can be compiled separately as a "loadable module." For ease of use, the PPC uses both methods. Drivers compiled into the kernel are always active, and require no user or system input to work. Drivers such as TCP/IP, SLIP, PPP, ramdisk, and serial have been compiled into the kernel.

Modular drivers must be explicitly loaded into the kernel, after the machine has booted. They can be loaded (and unloaded) at any time. Most commonly, they are loaded by boot scripts (similar to the DOS autoexec.bat script). Drivers such as parallel port, DOS filesystem, and ethernet have been compiled as modules. Modules are kept in subdirectories of the /lib/modules/2.2.14/ directory, indexed by type. Modular drivers are loaded by the insmod command, and removed by the rmmod command. All modules currently loaded can be listed by the lsmod command. Note that some drivers require parameters on the command line. In particular, ISA peripherals like network and sound hardware usually need their I/O base address specified.

The PPC automatically loads a number of modules at boot time. Any module listed in the file /etc/modules will be loaded at boot time. Command line parameters for these modules may also be specified. Note that for space reasons, the PPC includes very few of the available modules. If you require a driver which is not included by default, contact EMAC to have the driver sent to you.

Another caveat of using modules is that some modules depend on others, and will not load until the dependencies have been satisfied. The PPC cannot automatically resolve these dependencies, so a user must explicitly specify the loading of any module dependencies before loading the desired modules. Some common examples of dependencies are the ne2000 driver (ne.o), which requires the 8390.o module, and the SLIP and PPP drivers (slip.o and ppp.o), which require the slhc.o module.

Section 4. Flash disks and RAM disks

The PPC runs its operating system from an M-Systems DiskOnChip. This is in contrast to some other small Linux distributions, which load the filesystem into a RAM disk and execute from there.

However, the PPC does use RAM disks for some tasks. Flash media such as the DiskOnChip can only support a limited number of write operations (though they can be read an unlimited number of times.) As such, a RAM disk is used for the directory tree which is most often written to: the /var directories.

The /var directories contain files which are intended to be changed often. Lock files, temporary files, and log files are all kept in /var. Keeping the /var directories in a RAM disk saves a great deal of wear on the flash. However, it also means that log files are lost any time the PPC loses power. If you wish to keep log files on a permanent basis, you should use the periodic command scheduler CRON (described later in this document) to copy the logs to flash, or to send them to a remote computer via FTP or e-mail.

As with most other operating systems, a Linux system should generally not be powered off without first shutting down the operating system. This allows Linux to flush its buffers and finish any pending writes to disk (flash). The PPC can be properly shutdown by using the command line "shutdown -h now" (substituting an "r" for the "h" will cause the PPC to reboot, instead of halting and waiting for power off). If a Linux system is powered down without warning, data may be lost. In the worst case, a power down could occur while a write is taking place, causing corruption of the filesystem. This is particularly dangerous when using flash, as a flash device must be erased before it can be written, greatly increasing the time in which the media is vulnerable to corruption.

Read-only filesystem:

To work around the problem of filesystem corruption, it is possible to mount the filesystem read-only. If no writes can be performed, no corruption can occur. Mounting read-only is done in the boot scripts. If the file /etc/READONLY exists, the PPC will attempt to mount its root filesystem read-only. However, the user will be given the opportunity to override this each time the system boots. During the boot script sequence, a prompt will appear, and the user will be given five seconds to press a key to override the read-only option. If the key is not pressed, the filesystem will be mounted read-only.

It is also possible to change the filesystem from read-only to read-write and back from the command line. The command "mount -o remount,rw /dev/tffs1 /" will re-mount the filesystem in read-write mode. Using the same command line, but replacing "rw" with "ro", will remount the flash read-only.

Note that when the root filesystem is mounted read-only, the /dev directory is mounted in a second ramdisk. This allows users to log in even when the filesystem is read-only. If you use the command lines above to remount the flash read-write, changes to /dev will be lost when the system loses power.

Disk compression:

There are times when disk space is the premium concern in an embedded system. To help alleviate this problem, EMAC has included support for on-the-fly disk compression in the PPC's kernel. Depending on the type of data on the flash, using this compression can increase available flash space by up to 90%. This compression is entirely transparent to the user, requiring no special programs or additional setup.

It is important to realize that filesystem compression is not appropriate in all situations. When compression is in use, disk performance suffers greatly, depending on the processor. Compression also adds another potential point of fault in the system by adding another layer of complexity in the base operating system. Finally, standard GNU/Linux filesystem tools will not work with a filesystem using compression, making recovery that much more difficult in the case of an emergency.

There are two commands which affect the use of compression in the PPC filesystem. The lsattr command will list the contents of a directory, displaying their attributes, including compression status. The chattr command will change those attributes. Attributes can be changed on single files, or on entire directories. The command line "chattr +c <filename>" will compress a file, and the command line "chattr -

c <filename>” will reverse the compression. If a ‘c’ shows up in a file’s attributes as displayed by lsattr, the file is compressed. Compression can be controlled from the master configuration utility, described later in this document.

When a file has the compression flag set, it will be compressed immediately, and all subsequent writes to that file will likewise be compressed on-the-fly until the compression flag is unset. When a directory’s compression attribute is set, every file in that directory is immediately compressed. Furthermore, any file written to that directory in the future will also be compressed, until the directory’s compress attribute is unset.

IT IS VERY IMPORTANT THAT SOME DIRECTORIES AND FILES NOT BE COMPRESSED, SUCH AS THE KERNEL AND BOOTLOADER. COMPRESSION OF ANY FILE IN THE /boot DIRECTORY WILL RESULT IN A NON-FUNCTIONAL PPC. DO NOT SET THE COMPRESSION FLAG ON THE /boot DIRECTORY, OR ANY FILE IN IT. ADDITIONALLY, DO NOT COMPRESS THE / DIRECTORY RECURSIVELY (INCLUDING SUBDIRECTORIES,) AS THIS WILL COMPRESS THE BOOT DIRECTORY.

The first time a file on the PPC is compressed, a persistent flag is set in the filesystem’s master information area (“superblock”). This flag will NOT be unset when the last file on the PPC is uncompressed. While this flag is set, normal filesystem tools which are unaware of compression (such as those tools found on every major, modern Linux distribution) will be unable to mount, read, or otherwise deal with the PPC’s filesystem. To clear this flag, all files on the PPC must first be uncompressed. Then, the utility clear-e2c must be run. This utility’s sole purpose is to unset the “compressed” flag on the filesystem superblock, if appropriate.

Section 5. Login Terminals

There are (at least) four major ways to access the PPC. Each is largely similar to the others, but there are a few important differences. The four terminals described here will be the console, serial terminals, dial-in terminals, and TELNET sessions.

The first way to access the PPC is via the console. The console is the way most people think of using personal computers: type at a keyboard plugged into the unit, get feedback on a CRT, also plugged into the PPC. This is perhaps the easiest way to use the PPC, as no cables or configuration is required. The console also has the advantage that all system messages are displayed to it. These messages are generally not displayed to other terminal types. From the console, it is also possible to watch the system startup messages, and override certain system options (like mounting the filesystem read-only) at boot time.

In the PPC’s default configuration, a serial terminal server is attached to ttyS0 (COM 1). The serial settings for this terminal server are 9600 baud, 8 data bits, 1 stop bit, no parity, vt100 terminal type. The serial terminal settings are configured in the /etc/inittab file, and should be controlled through the master configuration utility (described later in this document).

When a “dumb terminal” is plugged into this serial port, a login prompt appears, and a user can log into the PPC. Alternately, a second PC can be plugged into ttyS0 using a NULL modem cable, and a terminal program such as minicom or hyperterminal may be used. It is important that when using a second PC, a “NULL modem” or “crossover” cable is used; a standard serial cable will not work.

The third type of terminal available on all PPCs is the TELNET server. For this to work, the PPC’s ethernet port must be plugged into a network, and its networking must be configured properly. In particular, you must set the PPC’s IP address and routing data to reflect your LAN. If you are using DHCP, you must know what IP address your DHCP server assigns the PPC. You can test the network setup of the PPC by using the “ping” program found on most network operating systems. If the PPC responds on the expected IP address, you can then use the TELNET client included with the operating system on another machine to TELNET to the PPC and log in.

Device-based security:

To increase security, the root user is only allowed to log in via certain "devices." Every terminal type uses a different type of device. The root user must be explicitly allowed to log in on each device. However, since UNIX/Linux allows for multiple users, there are multiple devices of each type.

The console uses devices tty0 through tty3. Serial terminals use devices ttyS0 - ttyS3 (these correspond to COM1 - COM4). Since modems are generally treated as serial devices, dial-in terminals also use a ttySx device, depending on which serial port the modem occupies (usually ttyS2). TELNET uses devices tty0 and so on...there are many tty* devices, as many users may be logged in via TELNET at once.

The root user may only log in on devices explicitly listed in the file /etc/security. Should changes need to be made, this file must be edited by hand. By default, several TELNET devices are listed in this file. This aids remote configuration and troubleshooting, but it is also often a security risk. If you do not intend to configure or manage the PPC via TELNET, it is recommended that the TELNET devices (tty*) be deleted from this file. IT IS IMPORTANT THAT THE CONSOLE AND SERIAL DEVICES NOT BE REMOVED FROM THIS FILE. If that happened, then root would be unable to log in, and configuration and management of the PPC would be impossible.

Login scripts:

When a user logs in, certain scripts are automatically run. These scripts will set the user's path, prompt, and other environmental variables. The scripts /etc/profile, ~/.profile, and ~/.bashrc are all executed, in order. The scripts in your home directory may be changed should you wish to add to the path, set a command alias, or to run a program every time you log in.

This is the case with the root user: every time the root user logs in, the master configuration utility is started by the ~root/.profile script. Note, however, that only the ~/.bashrc script is executed if a shell is not the user's login shell. Thus, the master configuration utility is not started again when the root user requests a second shell (bash prompt) from the master configuration utility.

Section 6. Periodic Command Scheduler

The PPC comes with the standard UNIX periodic command scheduler, CRON (Vixie implementation). This is a program which starts at boot time and runs in the background, executing specified commands at specified times. Programs can be executable binaries or shell scripts. CRON is configured by editing its configuration file /etc/crontab, which may be done from the main system configuration utility. The file may also be edited directly.

The crontab file consists of zero or more entries with the basic format "at this time, execute this command." Since CRON executes on its own (with no user account configuration), crontab may also include statements which set and export the PATH shell variable (the default crontab file contains such statements, which should not be modified except by an experienced user).

Times may be specified from the minute up to the month. The specific format for a crontab entry is:

<minute> <hour> <day of month> <month> <day of week> <user> <command>

All times may be specified by number, and the hour field goes by a 24-hour clock. The day of the week may be specified numerically, by the first three letters of its name. When specified numerically, the day of the week may be 0-7, with both 0 and 7 representing Sunday. Ranges and lists of times are supported. Thus, placing "5-10" in the hour field will cause the command to be executed every hour from

five AM to ten AM. The equivalent (comma-separated) list would be "5, 6, 7, 8, 9, 10". Note that lists and ranges are not supported when using week day names. Divisors are also supported. For example, "5-10/2" in the hour column would execute the given command once every other hour between five and ten AM. Finally, "any" or "all" may be specified by placing an asterisk (*) in the field.

For example, the line

```
42 0 * 1,5 2-4 root /bin/echo "CRON command is executing" > /dev/console
```

will print the message "CRON command is executing" to the PPC's console at 12:42 AM every Tuesday through Thursday during January and May. The field "user" contains the user that the command will be executed as. In most cases, this will be **root**. (Note that **root** has special privileges; care should be taken that essential system files are not modified by crontab entries.)

Crontab has several pre-defined entries, designed to perform standard system "housekeeping" tasks. These entries are designed to run every file (usually shell scripts) in a particular /etc subdirectory at regular intervals. These subdirectories are cron.daily/, cron.weekly/, and cron.monthly/. Simply put, everything in /etc/cron.daily/ will be executed once per day (usually late at night so as not to slow down user operations.) Everything in /etc/cron.weekly/ will be executed once per week, and so on. This can be a good mechanism to limit the lengths of log files, and basic scripts to manage log files are already in the CRON directories. User files placed into these subdirectories will be executed along with the log file management scripts, providing a simple alternative to making actual crontab entries.

Section 7. Developing Your Own Applications

The PPC is a real Linux system, and as such can run any application that can be run on a Linux workstation. However, the PPC is not a development platform itself. It has no compiler, no standard headers, no help tools, and a rudimentary text editor.

To develop custom applications for the PPC, it is strongly recommended that the user set up a Linux workstation using one of the modern, standard Linux distributions. RedHat 6.x, Debian 2.2, SuSE 6.x, and many others will work well. For maximum compatibility, EMAC also recommends that your development system use the same kernel and standard C libraries as the PPC (the PPC uses kernel 2.2.14 and glibc 2.1.3). In particular, we cannot guarantee compatibility if your applications is compiled on a system with a newer version of glibc. All executables should be compiled on this development workstation, then transferred to the PPC via FTP (remember to set permissions on the file to allow it to be executed!).

Many languages and compilers are available for Linux, including C/C++, Fortran, BASIC, assembly, and Perl. Not all languages are supported by default. For example, the default PPC configuration does not include interpreters for BASIC, Perl, Python, or any other interpreted language (shell script excepted). Compiled languages such as C, Fortran, and assembly do not require special interpreters to execute. However, a particular program may need certain "shared objects" added to the system before they will execute.

Shared libraries:

To save disk space, GNU/Linux systems keep code used by multiple applications in "shared object" libraries. The libraries work much like .dll files in MS Windows. The libraries are usually found in the /lib, /usr/lib, and /usr/local/lib directories, and can be identified by their .so.<version> extensions. The /lib directory also contains the dynamic linker, ld-linux.so.2.

Because flash space is limited, the PPC only contains the most commonly used libraries, such as the standard C libraries, network name resolution, and the ncurses library. Many programs will need other libraries to run. If your application will not run, use the command line "ldd <filename>" to display a

list of the libraries your application needs. If any of the libraries are listed as "not found," then you will need to either copy that library to the PPC, or specify static linking when you compile your program.

It is also important to note that many "libraries" are actually dynamic links to files with similar names. Depending on the compiler, an application may search for a library by several different names. If you copy a library to the PPC and your application still does not find it, you may need to create a link to the library with a different name. Shared objects are also cached, along with their locations. If your application continues to insist that a library does not exist even though it does, you may need to use a program called `ldconfig`. Copy this binary to the PPC from your development workstation, and execute it once, then try your application again -- odds are that it will now function properly.

Troubleshooting:

If your program works properly on your development workstation but reports "segmentation faults" when run on the PPC, there are several possibilities. The first is that your program was linked with a newer version of the standard C libraries, and the PPC's libraries are not forward-compatible. This behavior was seen when programs linked against `glibc 2.1.2` were run using `glibc 2.0.7`. To reduce this problem, the PPC was created using the newest libraries available at the time. However, technology is sure to progress.

The second major possibility is that your program binary was corrupted in transfer. In particular, the DOS-based FTP client included with MS Windows '95 and '98 default to ASCII mode, meaning that the FTP client may modify the file it is transferring. This FTP client must explicitly be set to transfer binary files. Other FTP clients may have similar behavior.

If you get a "No such file or directory" error message when attempting to execute your program on the PPC, there are two possibilities. First, as a security measure, the current directory is not in the search path. The command interpreter will only look for executables in the search path. In MS-DOS, the current directory was part of the search path; therefore executables in a user's current directory could be executed without prepending any directory information. To execute a file in the current directory, prepend a period and a slash (`./`) to explicitly specify execution from the current directory.

The second possibility is that there is another sort of library problem. The "No such file or directory" error can be caused when a program compiled and linked against `libc` version 5 is run on a system without this older library (such as the PPC). Unfortunately, the error message is rather uninformative. Use `ldd` on the binary to see if it requires `libc5`. If binaries you create do require `libc5`, EMAC recommends you update your development workstation to a newer distribution which uses the newer libraries, and recompile. Slackware in particular has been cautious and slow in adopting the newer standard C libraries. If it is absolutely impossible to recompile the application with newer libraries, the following files may be copied to the `/lib/` directory of the PPC: `libc.so.5`, `ld.so`, `ld-linux.so.1`. Each of these files must have their permissions set to 755. These files will require in excess of 750 kB of flash space. You may also want to use the `strip` utility (found on your development workstation) on these files before transferring them to the PPC, as this may cut the file sizes significantly.

Starting applications at boot time:

It can be desirable to have an application start at boot time. There are two ways to accomplish this: via setup scripts, or via `init`. `init` is that "mother of all processes." Its task is to start login terminals on specified devices, and usually keep them running by re-starting them when they exit. `init` is configured by the `/etc/inittab` file. `init` is designed to associate tasks with devices. If your application will use a single serial port or a virtual console terminal, then `init` may be an appropriate method to run your application each time the PPC boots.

If your application is not associated with a single device, then it should be started via boot scripts. These scripts perform much the same tasks as the DOS `autoexec.bat` script, though on a much grander scale. PPC boot scripts are stored in the `/etc/init.d` directory, however they are not executed from that

location. When init starts, it executes a script called rc.d. This script in turn starts several other scripts. First, scripts in /etc/rcS.d is executed, in alphabetical order. Only scripts that start with the letter 'S' are executed, and it is common to follow this 'S' with a two-digit number, to control order of execution. Note that these scripts in /etc/rcS.d are actually links to the actual scripts in /etc/init.d.

Once the scripts in /etc/rcS.d have been executed, another directory of scripts is executed. Which directory depends on the "run level" of the system. The run level is a concept used by UNIX to make it easy to control the starting and stopping of different services at boot time. The default run level is set in /etc/inittab. The PPC's default run level is 2. When run level 2 is entered, scripts in /etc/rc2.d are executed. Similarly, when run level 5 is entered, the scripts in /etc/rc5.d will be executed.

When the system leaves one runlevel, it first attempts to "kill" services which are no longer desired, by running scripts in the appropriate directory which begin with 'K'. For example, when the system leaves run level 2, all scripts in /etc/rc2.d beginning with 'K' are executed. The start scripts (scripts that begins with 'S') are then executed in the directory for the new run level.

To start you application at boot time, create a shell script which starts your program and place it in the /etc/init.d directory. Be certain to set execute permission on your script. Then, in the directory associated with the desired run level, create a link to your script. Make sure the link name begins with 'S', and use a number which falls after the rest of the scripts ('S99' is usually safe). Your script will then run after all other system configuration has taken place, and start your application.

Section 8. Networking addressing

The PPC communicates with other computers using the standard suite of internet protocols (IP). Computers using IP talk to one another by means of an IP address. Each machine on the Internet has a unique IP address, assigned by their network administrator. An IP address usually consists of four numbers (0-255), separated by dots. Examples are 192.168.1.50 and 10.0.2.1. By default the PPC uses DHCP and will automatically try to obtain its IP address from the network when using Ethernet.

The PPC has several ways in which it can connect to a network. The PPC can use IP over one or more serial ports (SLIP devices,) modems (PPP devices,) or ethernet devices. These devices are referred to as "interfaces". A particular interface is described by a combination of letters, which describe the type of interface, and a number. Interfaces of a particular type are numbered from 0, in the order that they were activated. For example, the first ethernet interface activated is identified as "eth0". Similarly, the first SLIP device activated is "sl0," and the first PPP device is "ppp0." Note that eth0, sl0, and ppp0 may all exist on the same machine at the same time.

An PPC is not restricted to having a single IP address. In fact, every network interface may have a different IP address. For example, the eth0 interface may have address 10.0.2.41 (the factory default). IP packets arriving at that interface addressed to 192.168.0.5 would be ignored. However, at the same time, the ppp0 interface may be assigned IP address 192.168.0.5, and accept packets so addressed.

Networks and routing:

Different IP addresses on the same PPC brings us to the concepts of networks and routing. The PPC's kernel holds an internal routing table, which tells the network subsystem where to find remote computers. Specifically, the routing table holds zero or more entries stating which computers or "networks" may be found on a particular interface. For example, a single routing entry may specify that packets addressed to the computer with IP address 124.172.18.9 should be sent out over interface sl0.

A "network" is a group of computers with similar IP addresses. A network of computers may have one or more sections of their IP addresses in common. For example, computers with IP addresses 10.0.2.4 and 10.0.2.5 are on the same network. The number of sections the network machines have in common is determined by the "netmask". When the IP addresses and netmask are converted to binary, the netmask contains 1's in all locations where the IP addresses match. Zeros in the netmask may not appear to the left of 1's. For example, 10.0.2.5 and 10.0.2.78 are in a network with netmask 255.255.255.0 (this is called a "class C" network). Machines with addresses 127.6.83.9 and 127.4.14.62 would be on a network with netmask 255.0.0.0 (a "class A" network, very uncommon).

Each network also has a "broadcast address". The broadcast address is an IP address consisting of all of the sections of the IP which the machines have in common, and 0's in the rest. For example, in a network with machines numbered 192.168.*.* and netmask 255.255.0.0, the broadcast address is 192.168.0.0. All machines on a network will receive a packet to the broadcast address.

Default gateway:

It is possible (in fact, common) for an IP packet to be addressed to a machine which cannot be located in the PPC kernel's routing table. These packets are dealt with by means of a "default gateway." An PPC may have zero or one default gateway(s). The default gateway must be locatable by the kernel's routing table. Whenever a packet is encountered which is addressed to an unknown machine, it is sent to the default gateway. This gateway is almost always a switch, bridge, or router between one or more networks.

Command-line network tools:

Each interface has an IP address, netmask, and broadcast address associated with it. This information may be set or viewed via the `ifconfig` command. With no parameters, `ifconfig` will display information for every active interface. Alternately, a particular interface may be specified on the command line, whose information will be displayed whether the interface is active or not (though the device must exist; the driver must be loaded). Interface configuration should be done through the master configuration utility (described later).

The kernel routing table may be changed or viewed using the `route` command. With no parameters, the `route` command will display the entire kernel routing table using hostnames instead of IP addresses where possible (hostnames are described in the next section). Each routing entry consists of a host or network address, a netmask, and an interface. The default gateway is an exception; this entry does not include an interface. The kernel must be able to locate the default gateway using the rest of the information in the routing table. Routing information should be changed through the master configuration script.

If your PPC does not use nameservers (described in the next section,) it is recommended that you use the `-n` command line switch with `route`. If you do not, the program will hang while it attempts to resolve IP addresses into hostnames, which will ultimately fail.

DHCP:

Some networks want to assign IP addresses and other interface information using a protocol named DHCP. The PPC includes a DHCP client, which is capable of automatically configuring the `eth0` interface, as well as adding routes and nameservers (described in the next section). You must have a DHCP server configured on your network to use DHCP. Use of DHCP should be specified when configuring the `eth0` interface with the master configuration script.

Section 9. Name resolution

Internet servers are frequently referred to by a "hostname", rather than by IP address. The hostname takes the form of two or more strings of characters separated by dots, such as `www.emacinc.com`. The sections of the hostname generally have no correspondence to the sections of the IP address.

Internet-wide hostnames are registered with one of a few companies, usually for a fee. A registrant must also provide the IP address that a hostname is associated with. Hostname / IP address pairs are stored on DNS machines (Domain Name Service). DNS servers may run on machines with other services, such as mail and web server. Your company network and/or ISP should be able to provide you with a list of name servers. Note that name servers are always referred to by IP address; the paradox is obvious.

A hostname does not need to be internet-wide; every UNIX box has and uses a hostname, at least internally. The factory default hostname for the PPC is `"PPCx20.emacinc.com"` (no quotes.) The hostname of the PPC is set at boot time, using the `hostname` utility. Without no arguments, this utility will also display the system's current hostname. The name set at boot time is the name contained in the `/etc/hostname` file. The contents of this file may be edited manually, or through the master configuration utility. Note that simply editing this file does not change the hostname of a running system; the `hostname` utility must be used, or the system may be rebooted.

The PPC has two ways to resolve hostnames into IP addresses. The simplest is locally, via the `/etc/hosts` file. This file contains a local listing of hostname / IP address pairs which are accessible only to the PPC. This file always includes one entry, the `localhost / 127.0.0.1` pair. Generally, this file should also include a second entry relating the PPC's set hostname with the IP address of interface `eth0`. This

entry is important, as some network applications will refuse to run if they cannot resolve the hostname of the PPC into the IP address of the interface they wish to use.

The second way the PPC can resolve hostnames into IP addresses is via DNS servers. Valid DNS servers should be listed in the `/etc/resolv.conf` file. Two to three DNS nameservers are common, but more are allowed. Note that if an application on the PPC attempts to resolve a hostname to an IP address via DNS servers, the application may stop for several minutes while the PPC attempts to contact the DNS server, especially if the DNS is inaccessible and the operation times out.

By default, the `/etc/hosts` file is searched before a DNS is accessed to resolve a name. This behavior may be altered by the `/etc/host.conf` file, which must be edited by hand. However, it is recommended that only an experienced Linux user do so. Note also that some DHCP servers will assign DNS servers to their clients. The PPC's DHCP client is capable of receiving and this data, and storing it in `/etc/resolv.conf`. However, this is not possible if the root filesystem is mounted read-only.

Section 10. Access control

Many common internet services (HTTP, FTP, TELNET) on the PPC are managed through a "superserver" called `inetd`. This superserver reads all incoming IP traffic, then passes the data to the appropriate service. Since all IP traffic passes through this single program, this allows for unified access control; the same configuration files can control access to many different internet services.

Control begins with the `/etc/services` files. This file contains a master list of all TCP and UDP ports that `inetd` will listen on, and the appropriate service to pass data on that port to. If a port is not listed in this file, then `inetd` will ignore traffic on it (though other stand-alone network servers may listen on that port.) In general, this file should not be edited unless security is a major concern, and then only by those already knowledgeable about Linux.

The next line of defense are the files `/etc/hosts.allow` and `/etc/hosts.deny`. As their names suggests, these files allow a system administrator to explicitly list servers, services, users, networks, and domains which may or may not use the `inetd` services (a "domain" is the hostname equivalent of a network).

The `hosts.allow` file takes the greatest precedence. If an entity is granted access by the `hosts.allow` file, then access will be granted to that entity no matter what is listed in `hosts.deny`. The `hosts.deny` file lists entities which should be denied access to the PPC. If an entity appears in neither file, the default action is to allow access. By default, the PPC ships with a relatively permissive access policy. `Hosts.allow` and `hosts.deny` may be edited by hand should you wish to alter the access policy. EMAC recommends this be done only by an experienced Linux user. The format for these files is beyond the scope of this document.

The FTP server is configured by the files `/etc/ftpaccess`, `/etc/ftpusers`, `/etc/ftphosts`, and `/etc/ftpgroups`. Not all of these files are present on the PPC, though they may be added if desired. `Ftpaccess` is the main FTPD configuration file, the others are used for fine-grained access control. The format of these files is beyond the scope of this document. Configuration of the FTP server is accomplished by editing these files by hand. EMAC recommends this be done only by an experienced user.

Section 12. PPP dial-in and -out

The same program works as both PPP server and client on the PPC. The `pppd` program may be invoked in one of two ways, both of which are simplified by scripts. The PPP protocol was designed to use dynamic IP addresses. This means that the PPP client (usually the machine doing the dialing) gets

its IP address and interface configuration from the PPP server (usually the machine being dialed) before any IP packets are exchanged.

Dial-out / PPP client:

To use the PPC as a PPP client, you must first create a dial-out script, then execute it. Dial-out scripts should be created with the master configuration utility. These scripts are stored in the `/etc/ppp/` directory. These scripts can be executed by typing the name of the script on the command line, just like any other program. These scripts may also be invoked by the CRON program, so that the PPC can connect to the internet via phone at specified times and upload collected data to a central server, for example.

PPP dial-out scripts will dial a phone number, enter a username and password, then optionally enter a command to start ppp on a remote machine. If further interaction with a remote server is required, an experienced Linux user may edit the `/etc/ppp/ppp-on-dialer` script.

Dial-in / PPP server:

The script which invokes pppd as a server is `/usr/bin/ppp`. This script will start the PPP server on any login terminal, including the console and TELNET sessions. However, it is generally only useful when started from a dial-in terminal. To use the PPP server in the normal sense (as you would dial your ISP from home), you must first set the PPC's modem to answer the phone as a dial-in terminal. Once the dial-in terminal connection has been established, you must log in as a non-root user (usually `www`). After logging in, a PPP session may be started by simply typing "ppp" (no quotes) on the command line. This process of logging in and starting PPP is usually automated by dial scripts on the client machine. Once PPP has been started, unreadable character strings will begin to appear in the dial-in terminal. This is PPP data; at this point you should instruct the client machine to begin a PPP session.

When acting as a PPP server, the PPC needs two IP addresses assigned to a ppp interface before it is started: a local address which the PPC will use for the ppp interface, and a remote address to assign to the calling, client machine. For simplicity and logical reasons, these IP addresses are associated with serial devices rather than with network interfaces. These IP addresses are stored as `/etc/ppp/options.ttyS*`. Since the PPC's modem is on `ttyS2` by default, IP addresses will be stored as `options.ttyS2`. The format for these addresses is `<local IP>:<remote IP>`. An IP address of `0.0.0.0` means that PPP should try to get the address from the remote host, rather than assigning an IP address to the remote host.

Since these IP addresses are assigned to serial devices, and since the program works as both PPP server and client, pppd will use the addresses whether it is acting as a server or a client. That is, if `options.ttyS*` specifies IP addresses, then pppd will attempt to use those addresses even if it is dialing out. This can lead to problems if the remote computer is not expecting to have its IP address assigned to it. While this can be difficult to manage if a modem will be used for both PPP dial-in and dial-out, it also allows for a very useful feature: dial-on-demand.

Dial-on-demand:

Dial-on-demand is a feature whereby a ppp interface may be configured and activated, without dialing the phone. The interface will sit idle until the PPC's network subsystem is given a packet which must be transmitted to the remote host on the ppp interface which is using dial-on-demand. The PPC will then dial the number, make the connection, and transmit the data. After a certain amount of time without PPP traffic, the PPC will hang up the phone and again wait for outgoing PPP traffic.

To use PPP dial-on-demand, two things must be satisfied: first, a dial-out script must be created. Second, the modem device must have at least a remote IP address assigned to it. If the PPC does not know the remote address on the PPP link, it cannot know when to activate the link by dialing the modem. The local IP address may be specified, or it may be retrieved from the remote host after dialing. After

creating the dial-out script in the master configuration utility, you will be given the option to make that script/link dial-on-demand. When this is done, the dial-out script will be run at boot time, and the ppp interface will be brought up. However, the phone will not dial until IP traffic addressed to the remote ppp IP address is encountered.

Section 12. HTTP/Web server

The PPC's web server is lightweight, single-tasking HTTP server named Boa (`/usr/sbin/boa`). In testing, Boa has been shown to be quite fast, able to serve hundreds of requests per second on a 486-class machine. Since it is single-tasking, Boa also uses much less memory than some more standard HTTP servers (such as Apache). Boa also takes comparatively little flash space. The trade-off is that Boa does not have many of Apache's advanced features, such as server-side scripting, ultra-fine-grained access control, or encryption. However, Boa has proven sufficient for most PPC tasks.

Boa is started at boot time by default. The script `/etc/init.d/netstd_init` starts both Boa (`http`) and the FTP server (`ftpd`). If this is not desired, it may be changed in two ways. First, the link `S03netstd_init` may be deleted from the `/etc/rc2.d` directory. This will prevent the `netstd_init` script from being executed when runlevel 2 (the default) is entered. Second, you may edit the `/etc/init.d/netstd_init` script by hand. Placing a `"#"` (no quotes) in front of the line which starts that `http` server will prevent the server from being started.

Configuration:

Boa's configuration is read from the `/etc/boa/boa.conf` file. This file contains directives for TCP port, user and group to run as, logging, and other configurations. An HTTP server usually runs on port 80; it is not recommended that this be changed. Once the server starts, it runs as user **nobody** and group **nogroup**. This is done for security reasons: the user `nobody` and group `nogroup` have no special permissions anywhere on the system. Therefore the HTTP server cannot read files that a non-privileged user should not, nor can it change file (via CGI, perhaps) that it should not.

To save RAM disk space, no logging is done. If you are having problems with the web server, logging may be enabled for both errors and HTTP accesses. The configuration file also allows you to set certain paths, directories, and file types. File types (MIME types) may be associated with file extensions both in the HTTP configuration file, and in the `/etc/mime.types` file. The paths tell the HTTP server where to find CGI and other files. In particular, Boa uses a CGI `/usr/lib/cgi-bin/boa_indexer`. This CGI generates directory listings on the fly when no `index.html` file is available. The `boa.conf` file is also somewhat self-documenting; see the comments in the file for more information.

One thing which cannot be set by the configuration file is the "server root." The server root is the directory which Boa considers the top-level directory; no higher levels can be accessed. For example, the default server root (compiled into Boa) is `/home/www`. When a browser is given the URL `http://10.0.2.41/` (assuming that the PPC is still at the default address 10.0.2.41,) the contents of the `/home/www` directory will be displayed – there is no mechanism by which the `/home` or `/` directories can be displayed. Note, however, that Boa will follow directory links where ever they may lead, which could potentially include higher level directories than the server root. A non-default server root must be set on the command line; this can be done via the `/etc/init.d/httpd` script.

Note that Boa is a stand-alone server, which does not use the `inetd`. Therefore, `inetd` access control may not work when using Boa. Definitive information on Boa may be found at www.boa.org (as of this writing).

Section 13. FTP server

The PPC's FTP server is the Linux-standard WU FTPD, version 2.4.2. This server is started by default at boot by the same scripts which begin the HTTP server, and it may be disabled in the same manner. WU-FTPD (`/usr/sbin/in.ftpd`) is run through `inetd`, thus all `inetd` access control methods will work as previously described. The file `/etc/ftpaccess` is the master `ftpd` configuration file. The files `/etc/ftphosts`, `/etc/ftpusers`, and `/etc/ftpgroups` all configure access control for the FTP server. The syntax of these files remains beyond the scope of this document; see the manual pages for these files on your development workstation for more detailed information. The PPC's FTP server uses the default ports 21 (for commands) and 20 (for data).

Section 14. TELNET server

The PPC's TELNET server (`/usr/sbin/in.telnetd`) is also the Linux standard. The TELNET server is run through `inetd`, thus all access control methods previously described will work. Further access control is accomplished by the same methods as limiting terminal access to any user – the username / password pair, as well as the `/etc/securetty` file. See the manual pages on your development workstation for more details. The PPC's TELNET server runs on the default TELNET port 23.

Section 15. Sending e-mail

To save space, the PPC uses a very small, simple program to transmit e-mail (reception of e-mail is not possible with the base PPC). The mail client is not a full-fledged mail transport agent; the PPC requires access to a properly configured SMTP mail server / transport agent. The setting up of such a machine is beyond the scope of this document. Your corporate network or local ISP probably have such a server, ask your system administrator for details. Usually, any machine on which you have an e-mail account will work.

The program to transmit mail is called "smtpclient." For a full list of options, type "smtpclient -help" on the command line. Smtplib is a command-line based program, and as such all options such as subject and destination must be specified on the command line. At minimum, you will need to use the "-S" or "smtp-host=" option to specify the IP address of the SMTP server, and a destination address (which does not require a command line switch). Depending on your mail server, the "-f" switch may also need to be used to specify who the message is from (some SMTP servers will not accept mail from users they do not recognize).

The smtpclient program reads its input from the "standard input." This means that you cannot use a switch to tell smtpclient the name of a file to send. The file must be passed by means of a "pipe" or by a "redirect." A pipe sends the output of one command to another. A pipe is specified by a "|" character. Use the cat command to print the desired file to the standard output, then use a "|" to pipe that output to the standard input of smtpclient: "cat /tmp/mailfile | smtpclient -S mail.your.com -s Mail from PPC -f you@your.com destuser@dest.com".

A redirect does something similar, but the file is specified after the smtpclient command line, and you don't need to cat the file. The command line "smtpclient -S mail.your.com -s Mail from PPC -f you@your.com destuser@dest.com < /tmp/mailfile" will do exactly the same thing as the command line using a pipe. Which you choose makes no difference. Mailing can be done from the command line, by a CRON task, or by a script. You may also write your own shell script to automate the mailing process.

Part II: Configuration

Getting started

The text-based PPC configuration program can be reached in one of three ways, either by telnet, the console, or by a serial terminal. To use telnet, plug the PPC ethernet interface into an ethernet port on a standard PC using a special point-to-point ethernet cable. Alternatively, the PPC and a standard PC could be plugged into an ethernet hub with no other devices attached.

The PPC is factory-configured to attempt DHCP configuration of its primary ethernet interface. If your network has a DHCP server, simply find out which address has been assigned to the PPC and TELNET to it. If The PPC fails to find a valid DHCP server, it will fall back to the factory defaults: IP address 10.0.2.41, netmask 255.255.255.0, broadcast 10.0.2.0, default gateway 10.0.2.1 on interface eth0. (For future reference, these defaults are stored in /etc/sysconfig/ethdefaults.)

If this is the case, you will need to configure your PC to communicate with the PPC directly. Under Windows 95/98, this means changing your default gateway to 10.0.2.41 (the PPC's factory-set temporary IP address). Under GNU/Linux, this means adding a routing entry for 10.0.2.41 on the appropriate device (via the "route" command). You may then telnet to the PPC at address 10.0.2.41 using your standard telnet program.

To log in via serial terminal, a "dumb terminal" may be attached to the PPC's primary serial port (COM1 or ttyS0). A standard PC running a terminal emulator such as Hyperterminal will also work. The terminal settings are 9600 bps, 8 data bits, 1 stop bit, no parity, no flow control, vt100 terminal settings. Just plug in the terminal, power on the PPC, wait a minute or so for the PPC to complete its boot process, then press enter a few times to "wake up" the terminal and bring up a login prompt.

The console option may be used if the video option has been installed. The console is what might be considered the usual way to use a computer: plug a keyboard and VGA monitor into the appropriate connectors, turn the PPC on, and wait for a login prompt.

Once you are connected to the PPC, log in as root. The default password is "emac_inc" (no quotes). As soon as the password is accepted, the main configuration menu should appear. EMAC recommends this password should be changed almost immediately from the main configuration menu. Note that if the flash disk has been mounted read-only, it will be re-mounted read-write while the configuration utility is running. it will then be mounted read-only again when the user exits from the program.

Main configuration menu

The main menu has seven main options. 'Serial settings' allows the user to configure a serial port as an IP interface (SLIP or PPP), a terminal interface (direct or dial-in), or a raw serial port. 'Ethernet settings' sets up the information for the PPC's primary ethernet interface, including IP address. The 'Routing tables' selection allows the user to add, delete, and view static routing information (i.e. which interfaces to use for particular destinations), including default gateway. 'Host name' allows the user to change the PPC's hostname. 'Internet name resolution' brings up the menu to configure DNS servers and manage static host entries. 'Configure periodic command scheduler' will bring up the contab configuration file in an editor, allowing CRON to be configured. 'Configure boot-time parameters' will allow you to configure what modules are loaded at boot time, and whether the flash is mounted read-only or read-write.

There are also four additional options. 'Bash prompt' will start an interactive shell (command prompt) from which commands may be run. 'Change root password' allows the user to change the access password to the PPC (for security purposes, it is recommended that this be done first.) 'Shutdown PPC' allows you to reboot or halt the PPC. Finally, 'quit' exits the program and bring up another login prompt. Each option is detailed here. Menu items marked with a "*" are those items necessary for almost every PPC configuration.

Serial settings*

Selecting 'Serial settings' from the main menu will bring up a submenu, containing seven selections plus an option to return to the main menu. These seven options are:

Configure SLIP port

"Configure SLIP port" configures a specified serial device as an IP network interface, with its own IP address. Note that this IP address will only be used for this particular SLIP port; every IP network interface may have its own different IP address if desired. To configure a SLIP port, you first must choose a serial device which will be connected to a network. In Linux, these devices are called ttyS0 - ttyS3 (sometimes with "/dev/" prepended,) which correspond to DOS ports COM1 - COM4. You must select the device by number alone. Note that you are allowed to configure serial devices which do not actually exist in hardware -- this will not cause any system errors or stability issues, but errors will occur later if data is written to or read from these non-existent ports. Make sure you specify a valid device. After specification of a device, a screen will appear which contains the present SLIP configuration for the given port. If there is no information, that means that the port has not been previously configured for SLIP.

To complete configuration of a given SLIP device, you will need to specify at least a line speed and an IP address. If the SLIP device will be connected to only one other computer/controller, 'Point to Point' should be specified. If the SLIP port will be connected to a network of other computers/controllers, do not specify point-to-point operation. If point-to-point is specified, you will be asked for the IP address of the machine to which the SLIP port will be attached. These IP addresses should be numeric address, not hostnames. After you have entered the data, the new configuration will be displayed, and you will be asked if you wish to proceed. If you type 'y', the device will be configured as shown (and all other configuration for that device will be deleted.) If you type 'n', no configuration data will be changed, and the device will remain configured as originally shown. Whether the data is modified or not, you will automatically be returned to the serial configuration menu. SLIP configuration data is stored in files named serial*ip in the /etc/sysconfig directory. These files are read and used by the /etc/sysconfig/serialcfg.sh script, which is run at boot time.

If multiple SLIP ports are configured, they should not share the same local IP address. For example, if both ttyS0 and ttyS1 are configured for SLIP, they should not both be given local IP address 192.168.1.77. This will cause confusion in the kernel routing tables. However, so long as each SLIP interface is allocated a separate IP address, as many SLIP interfaces as desired may be used.

Note that serial network interfaces are given names/numbers by the Linux kernel during system boot, in the order that the devices are configured. The first SLIP device configured is named "sl0," the second is "sl1," and so on. It is important to realize that these numbers are not necessarily equal to the serial port ttySx device numbers; /dev/ttyS3 could be sl0, and so on. However, the ttySx devices are configured in numeric order; ttyS0 is configured first, then ttyS1, and so on. The lowest numbered serial device configured as a SLIP port will be sl0, and so on. Thus if ttyS0 is configured as a terminal port, ttyS1 is a SLIP interface, ttyS2 is a raw serial port, and ttyS3 is a SLIP interface, ttyS1 will be sl0 and ttyS3 will be sl1.

It should also be noted that in the serial section of the configuration program, serial devices are referred to by their device numbers (ttySx), not their interface numbers (slx). However, under the routing submenu (detailed later), SLIP ports are referred to by interface number.

Configure terminal port

'Configure terminal port' allows the user to set up a terminal server on the specified serial port. Serial devices are specified in the same manner as they were in 'Configuring SLIP ports' (as ttySx.) To set up a terminal server on a serial port, you will be asked for the line speed and the type of terminal to expect on the port. The terminal type typically starts with the letters 'vt', common examples are "vt100" and "vt320". The "vt100" terminal type is the default, and is perhaps the most common terminal setting. After entering the new data, you will be shown the new configuration and asked for confirmation before any configuration data is changed. Again, accepting the configuration changes will delete any other configuration for the specified port. You will then be automatically returned to the serial configuration menu. Configuration data for all terminals is stored in /etc/inittab.

Configure dial-in terminal port

This menu item will allow the user to set up a terminal server which is modem-aware on the specified serial port. A serial device is specified as with previous menu items. When a serial port is configured this way and a modem is attached to it (or if the port is an internal modem), the PPC will answer the phone when it rings and attempt to make a data connection with the caller. If a data connection is made, the PPC will then prompt the caller for a login and password (in plain text), and the caller may log in as if they were using a standard terminal or the console.

The only parameter which must be specified is the line speed. As most modern modems are capable of high speed communications with the PPC serial port, this speed is 38400 by default. If the modem is internal, this parameter will be ignored by the hardware (though it is required as a placeholder for the software.) You will then be prompted for confirmation and returned to the serial configuration menu. Configuration for serial terminals, dial-in and standard, is kept in the /etc/inittab file.

Configure dial-in PPP port

This item allows a user to set up a dial-in port to allow PPP connections from a remote client computer. Note that to use a port for dial-in PPP, that port must first be configured as a dial-in terminal port, as described previously. Three items are required to complete this setup: the serial tty (ttySx) number where the modem is attached, a local IP address (the address the PPC will use for this PPP link), and a remote IP address (the address which the PPC PPP server will attempt to assign to the calling PPP client). If the calling client is to assign either or both IP addresses, then those addresses should be entered as 0.0.0.0. This will tell the PPP server that the client is going to assign that particular address remotely.

To start a PPP link, dial into the PPC and log in as the **www** user. When presented with a command prompt, simply type "ppp" (no quotes). This will begin the PPP connection. Note that most PPP-aware dialing software attempts to automate this login process. Check the documentation for your client software for more information.

Note that configuration data may exist for a dial-in PPP port, and may be shown in "View serial settings," even after that port has been re-configured as something other than a dial-in terminal port. This is not an error; dial-in PPP configuration data will simply be ignored for any port which is not configured as a dial-in terminal port.

Configure dial-out PPP connection

This menu item will automatically generate a script which, when executed, will cause the PPC's modem (on a specified serial port) to dial and attempt to make a PPP connection with a remote host. However, instead of being referenced by device number as in other serial configurations, these dialer scripts are referenced by a free-form identifier string. This string is chosen and supplied by the user when creating the script. This allows the user to create and store multiple dialing configurations for the same modem on the same serial device.

Upon selection of this menu item, you will be prompted for an identifier string, a phone number, a username (on the remote machine), the password for that username, and an optional netmask. You will also be prompted for the command prompt on the remote machine, and the command the remote host requires to start PPP. These last two items are only necessary if you are required to start PPP manually on the remote machine. If the remote host automatically starts PPP upon login, simply leave these answers blank. If these are specified, the dialer program will wait until it receives the text you specify as the remote command prompt, then send the string you specify as the command to begin PPP. Note that if you specify punctuation or other special characters, they may not be displayed properly in the configuration program, and may cause trouble with the dialer script. This trouble can be avoided by entering an alphanumeric-only substring of the remote command prompt.

Next, you will be asked if this dialer script is to be use dial-on-demand. If you say yes to this option, the phone will not dial when the script is run. The ppp* interface will be created, but the phone will not be dialed until there is IP traffic present addressed to the remote machine. Note that the dial-on-demand setting is specified for a serial device, rather than for a particular connection script. As such, setting dial-on-demand in one script will set it for all dial-out connections which use the same serial device.

Whether or not you wish to use dial-on-demand, you will next be prompted for local and remote IP addresses. If you want the PPC to try to retrieve these address from the remote host, leave both as 0.0.0.0. Remember that if you specify dial-on-demand, you must specify the remote host's IP address. Remember also that these IP address settings will affect all PPP connections on the specified device, both dial-out and dial-in.

Finally, you will be asked if you wish to activate this interface at boot time. This option is available for both dial-on-demand and standard connections. If you specify the interface to be activated at boot time, then your dial script will be added to the /etc/sysconfig/pSIBfg.sh script, and run at boot time. Keep in mind that this means the modem will dial at boot time, unless the script specifies dial-on-demand.

You will then be asked for confirmation before the script is generated. If you confirm this operation, a script with the name ppp-<identifier> will be generated in the /etc/ppp directory. Running this script will initiate a PPP link as specified. The general purpose ppp-off script may be used to terminate the PPP connection. Note that if there is more than one active PPP link (on more than one modem), you will need to specify the network device number (ppp0, ppp1, etc.) as a command line parameter to terminate the appropriate link; ppp0 is the default. Network device numbers for PPP devices are assigned in the order that PPP connections are made; the first link to dial is ppp0, the second is ppp1, and so on.

Remember, if you use dial-out PPP on a line which has a configuration file for dial-in PPP, the PPP daemon will attempt to use the dial-in IP data associated with the serial device when dialing out. For example, if you configured ttyS3 with the PPP dial-in data 192.168.0.1:192.168.0.70 (local IP:remote IP), then generated a dial-out script for ttyS3, the PPP daemon would use 192.168.0.1 as the local IP address for the dial-out connection, and attempt to assign 192.168.0.70 to the machine being called during dial-out. If this is not desired, re-configure the dial-in data with the addresses 0.0.0.0:0.0.0.0. This will tell the PPP daemon to

accept IP addresses from the machine being called (the PPP daemon will also accept IP addresses from the calling machine when answering a PPP dial-in call on the specified ttySx.)

*Configure standard serial port**

'Configure standard serial port' will allow the user to set parameters such as line speed, number of data bits, and parity of a raw serial device at boot time. To configure a raw serial port, you will be prompted for a line speed, number of data bits, number of stop bits, parity type, and type of flow control. Line speed may be any standard line speed from 1200 to 38400 inclusive (note that high-speed UARTs often use their maximum speed when 38400 is specified; this may be 57k or 115k for newer UARTs). Data bits may be from 5 to 8, inclusive. Stop bits may be 1 or 2. Parity may be even, odd, mark, space, or none, specified by E, O, M, S, and N respectively. 'Mark' and 'Space' parity are for 9-bit (multidrop) networks. Flow control may be Xon/Xoff software, RTS/CTS hardware, RS-422, or none, specified by s, h, 4, and n" respectively. If you use hardware flow control, make sure your cable has the appropriate number of wires. Note that some or all of these parameters may be overridden by user applications, including the gateway-cgi program.

If RS-422 is specified, then the RS-422 transmitter will be enabled at system boot, and will remain enabled during the entire operation of the PPC. Note that with standard PPC hardware, RS-422 is only available on port ttyS1 (COM2). Make certain your RS-422 network is connected to the correct hardware connector before using this option.

You will be prompted for confirmation before any changes are made, then returned to the serial configuration menu. As with other configuration options, accepting changes will delete any SLIP or terminal configuration for the specified serial device. This data is stored in files named serial*raw, in the /etc/sysconfig directory. These files are read and used by the /etc/sysconfig/serialcfg.sh script at boot time.

View serial settings

'View serial settings' will allow the user to view all settings for a particular serial port. This function will check for terminal, SLIP, dial-up terminal, and raw configuration, as well as PPP dial-in configuration and any PPP dial-out scripts which use the specified port. Any existing configurations will be displayed one item at a time. Note that if there are multiple configurations, this is most likely an error, with the exception of PPP dial-in configurations, which will be ignored if the device is not configured as a dial-in terminal, and PPP dial-out scripts, which will only be used if explicitly executed. (Note, however, that PPP dial-out scripts should not be explicitly executed if the port is configured as something other than a raw serial device.)

Ethernet Settings*

The ethernet settings menu has only two items, 'configure ethernet device' and 'view network device settings.' Ethernet devices are assigned interface names/numbers by the Linux kernel at boot time, in the order in which they are configured by the boot process. The first ethernet device configured will be named "eth0," the second will be "eth1," and so on. Ethernet devices are referenced throughout the configuration program by these interface numbers.

*Configure ethernet device**

'Configure ethernet device' will prompt the user for a device (number only,) then bring up a screen with the current interface configuration data. The user will then be prompted you for new data. To complete configuration of an ethernet interface, you will need an IP address (numeric, in form A.B.C.D), a netmask (frequently 255.255.255.0,) a broadcast address (frequently A.B.C.255,) and a network address (frequently A.B.C.0). After entering these parameters, you will

be prompted for confirmation and returned to the main menu. If you want an interface to use DHCP, then enter "DHCP" in place of the IP address. Ethernet configuration data is stored in files named eth*ip, and are stored in the /etc/sysconfig directory. These files are read and used by the /etc/sysconfig/ethcfg.sh script, which is run at boot time.

If DHCP configuration of an ethernet device fails for any reason at boot time, that interface will fall back to a set of defaults. At the factory, these defaults are configured to IP address 10.0.2.250, netmask 255.255.255.0, broadcast 10.0.2.0, and default gateway 10.0.2.1. These defaults are stored in the file /etc/sysconfig/ethdefaults, which is read by the ethcfg.sh script only after DHCP configuration fails.

View network device settings

'View network device setting' will display the configuration information for a given interface. As with serial configurations, no data means that the specified device is not a configured ethernet interface.

Routing tables*

This submenu will allow the user to modify the static routing table, which the PPC uses to determine which network device to use when trying to communicate with a particular host (the /etc/hosts file). The submenu items are 'Add routing entry', 'Delete routing entry', 'Set default gateway', and 'View routing data.'

*Add routing entry**

Under 'Add routing data,' you may add data for either a single host, or a network of hosts. At minimum, a PPC will need a host routing entry telling it on which network interface to find the default gateway (described in the next section). A single host will have a full address such as 10.0.2.1, and a netmask is optional but usually not required. You must then specify the network device the PPC should use to reach this host. This will be the name of a network device such as eth0 or sl1. Remember when attaching serial devices that the device (ttySx) number is not necessarily the interface (slx) number. However, boot initialization is set up such that the lowest number SLIP device will be given the interface number sl0. Thus if ttyS0 and ttyS1 are both raw serial devices and ttyS2 is a SLIP port, ttyS2 will be configured as sl0.

A network address will end in one or more zero fields, e.g. 192.168.0.0. This will tell the PPC that this entry is valid for all destination addresses with the form 192.168.x.x, where x can be anything from 1 to 254. For net entries, a netmask should be specified with zeros in the same locations as the zeros in the IP address, e.g. 255.255.0.0. You must then specify the network interface the PPC should use to reach this range of addresses.

Before any routing data is modified, the routing table entry will be displayed and you will be prompted for confirmation. The routing entry will be in the form:

```
/sbin/route add -host (or -net) <IP number> netmask <netmask> dev <interface>
```

If you did not enter a netmask, the section "netmask <netmask>" will not appear in the entry. Press 'y' to add the entry to the static routing table, or 'n' to abort any modification.

Note that PPP routing entries will be added automatically when a connection is made. Therefore it is not possible to configure a routing entry for a PPP device, as the kernel will reject this route until the interface exists (i.e. until the PPP daemon is running).

Delete routing entry

When 'Delete routing entry' is selected, a screen appears displaying the full current static routing table. This will consist of one or more entries of the form shown above, and one default routing entry of the form:

```
/sbin/route add default gw <IP address> metric 1
```

The meaning of the default route will be discussed in the following section. Below the routing table, you will be prompted for the IP address of the entry to delete. Note that entries will be matched by string comparison, so if you enter an address or other text common to more than one entry, then more than one entry will be selected for deletion. After you have entered the address to delete, you will be asked for confirmation, and any and all routes which would be deleted are displayed (this could be from none to all of the static routes, depending on the text you entered.) If you press 'y', all of the routes displayed under the confirmation request will be deleted. Note that if no routes appear following the confirmation request, then none will be deleted regardless of your answer to the confirmation request.

*Set default gateway**

The default gateway/route is where the PPC will send network packets when the destination of the packet has not been explicitly listed in the static routing table. The default gateway is generally your LAN's router, or other IP gatewaying/forwarding/switching device which will forward packets it receives onto another network. When you select the 'Set default gateway' option, a screen will appear which lists the current default route. If no route information appears, then no default gateway has been set. You will then be prompted for the (numeric) IP address of the new default gateway. This new default gateway entry will then be displayed and you will be prompted for confirmation. Entering a 'y' here will delete the previous default route and replace it with the one you have entered. Enter anything else and no changes will be made. Note that there should only be one default route in your routing table at any one time. Also note that prior to the definition of the default gateway, a static route to the gateway must be defined which includes a network device specification; if this entry does not exist, the networking subsystem will not know how to reach the default gateway.

View routing data

This option will display the full current routing table, including the default gateway. No modifications can be made from this screen. Pressing enter will return you to the routing table configuration submenu. Note that a standard host routing entry must be made for the default gateway prior to the specification of the default, or the PPC will not know how to reach the gateway.

Hostname

This option does not have a submenu. When selected, a screen will appear, stating the PPC's current hostname (this is pre-configured as SIBx20.emacinc.com, and should be changed to suit your network.) You will then be prompted as to whether or not you wish to change the hostname. If no, you will be returned to the main menu. If yes, you will be prompted to enter a new hostname, then prompted for confirmation before any changes are made. Note that unless this hostname is properly registered with your network administrator (possibly your internet service provider), this hostname will be meaningless except to the PPC itself. Also note that the PPC will not be able to resolve its own hostname into its own IP address unless an appropriate static host entry is made under the 'Internet name resolution' submenu (described below.)

Internet name resolution

When the PPC networking subsystem is given a hostname, it must be resolved to a numeric IP address before any data can be sent. There are two ways in which a name may be resolved into an IP address. The first method is through static host entries. The PPC maintains a list of hostname/IP address pairs (each of which also has an optional alias) in the file `/etc/hosts`. When given a hostname, the PPC first checks the hostname against this list, and if it finds a match, it uses the IP address found there. Static host entries may be managed by selecting 'Add static host entry' and 'Delete static host entry' from the name resolution submenu..

The second method by which a host name can be resolved into a numeric address is by communication with a name server. A name server is a computer which the PPC can access via network, which maintains an exhaustive list of host name / IP address entries. The PPC sends the host name it's looking for to the name server, and the name server sends a reply containing the numeric IP address of the host. Ask your network administrator or ISP for the IP addresses of any nameservers the PPC may have access to. The PPC stores its list of available nameservers in `/etc/resolv.conf`.

Add nameserver

When 'Add nameserver' is selected, the current name server information will be displayed. This will consist of a single entry in the form "search <domain>" and one or more entries of the form "nameserver <IP address>". It is common to have two or three name server entries. The <domain> in the search should contain the name of the domain in which the PPC will operate. If, for example, the EMAC computers were named `emac0.emacinc.com` through `emac10.emacinc.com`, the search entries would read "search `emacinc.com`".

Below the current nameserver information, you will be prompted as to whether you wish to change the search entry or add a nameserver (the default is add a nameserver). You will then be prompted for the domain name to search or the numeric IP address of a name server, respectively. The entry to be added to the name server configuration is then re-displayed, and you are asked for confirmation before changes are made. You will then be returned to the Internet name resolution submenu.

Delete nameserver

This submenu option will allow you to delete one or more entries from the current name server configuration. You will be prompted for the IP address of the entry to delete. As with the routing table, a string comparison is made, and as such multiple entries may be deleted at once (entering 'nameserver' will select all name server entries to be deleted.) After the address is entered, all lines to be deleted are displayed, and you are asked for confirmation before any modifications are made.

Add static host entry

This option allows you to associate a hostname with an IP address locally, so that the PPC will not need to contact a name server to resolve a host name. The full listing of static host entries will be displayed, and you will be prompted for a new hostname / IP address pair (with optional alias) to add. You should add entries for every remote host you wish to reference by name, but which do not have properly registered host names. You should also add entries for all hosts you wish to reference by name if you do not have access to a nameserver.

Delete static host entry

This option will allow you to delete an entry from the PPC's internal list of hostname / IP address pairs. The full listing will be displayed, and you will be prompted for the host name, IP

address, or alias of the host to be deleted from the list. As always, you are prompted for confirmation before any changes are made.

View Internet name resolution configuration

This option will display both the list of host names / IP addresses and the name server configuration. Pressing enter will return you to the internet name resolution submenu.

Configure periodic command scheduler

This option allows the user to directly edit CRON's configuration file, `/etc/crontab`. The crontab file consists of zero or more entries with the basic format "at this time, execute this command." Since CRON executes on its own (with no user account configuration), crontab may also include statements which set and export the PATH shell variable (the default crontab file contains such statements, which should not be modified except by an experienced user).

Times may be specified from the minute up to the month. The specific format for a crontab entry is:

<minute> <hour> <day of month> <month> <day of week> <user> <command>

All times may be specified by number, and the hour field goes by a 24-hour clock. The day of the week may be specified numerically, by the first three letters of its name. When specified numerically, the day of the week may be 0-7, with both 0 and 7 representing Sunday. Ranges and lists of times are supported. Thus, placing "5-10" in the hour field will cause the command to be executed every hour from five AM to ten AM. The equivalent (comma-separated) list would be "5, 6, 7, 8, 9, 10". Note that lists and ranges are not supported when using week day names. Divisors are also supported. For example, "5-10/2" in the hour column would execute the given command once every other hour between five and ten AM. Finally, "any" or "all" may be specified by placing an asterisk (*) in the field.

For example, the line

```
42 0 * 1,5 2-4 root /bin/echo "CRON command is executing" > /dev/console
```

Will print the message "CRON command is executing" to the PPC's console at 12:42 AM every Tuesday through Thursday during January and May. The field "user" contains the user that the command will be executed as. In most cases, this will be **root**. (Note that **root** has special privileges; care should be taken that essential system files are not modified by crontab entries.)

Crontab has several pre-defined entries, designed to perform standard system "housekeeping" tasks. These entries are designed to run every file (usually shell scripts) in a particular `/etc` subdirectory at regular intervals. These subdirectories are `cron.daily/`, `cron.weekly/`, and `cron.monthly/`. Simply put, everything in `/etc/cron.daily/` will be executed once per day (frequently late at night so as not to slow down user operations.) Everything in `/etc/cron.weekly/` will be executed once per week, and so on. This can be a good mechanism to limit the lengths of log files, and basic scripts to manage log files are already in the CRON directories. User files placed into these subdirectories will be executed along with the log file management scripts, providing a simple alternative to making actual crontab entries.

Configure boot-time parameters

This selection has three sub-options: 'Set root filesystem READ-ONLY,' 'Set root filesystem READ-WRITE,' and 'Configure loadable modules / drivers.' When the first item is selected, the file /etc/READONLY will be created, which will cause the flash disk to be mounted read-only the next time the PPC boots. If the filesystem is already set to be mounted read-only, this option will have no effect. Remember that even when the system is set to mount read-only, this may always be overridden at boot time (provided you have access to the PPC console).

When the second option, 'Set root filesystem READ-WRITE,' is selected, the /etc/READONLY file will be deleted, and the flash disk will be mounted read-write the next time the system boots. Read-write cannot be overridden at boot time. Note that neither of these two options have any immediate effect; they will only take effect when the system reboots. Also note that you will not be prompted for confirmation for either of these operations.

The third sub-option, 'Configure loadable modules / drivers,' will bring up the file /etc/modules in an editor. Any loadable kernel modules listed in this file will be loaded at boot time (before any network configuration takes place). Keep in mind that some modules will not load unless others are loaded first. The PPC cannot automatically resolve these dependencies -- it is up to the user to insure that modules and their dependencies are loaded in the correct order.

Bash prompt

This option will start a standard bash shell (command line interface), from which any command-line tools or scripts may be started and / or used. Type "exit" (no quotes) to leave the shell and re-enter the configuration utility. Remember that that filesystem is always mounted read-write while the configuration utility is running. Thus writes will always be enabled when you select this option.

Change root password*

Selecting this option will allow the user to change the password for user **root**. As this password is essential to system security, it is recommended that this password be kept as secret as possible, changed frequently, and never written down. The program which changes the password performs some basic checking to determine if an entered password is a "good" password; minimum length and certain content requirements must be met before a new password will be accepted. This makes passwords more difficult to acquire through direct decryption methods. As always, you will be asked for confirmation before anything is changed (though in this case, before any information is entered.)

The password for user **www** must be changed by logging in and running the program **passwd**. This will bring up the same verification and new password dialog as seen in the configuration program. The **root** password may also be changed in this manner.