/bin there, done that

Chris Bidmead throws light on Unix directories — /bin, /dev, etc. It's an informal layout, but you *can* get a feel for how it's put together. Plus, there's ping, SCO, and the great awk.

Last month I described how the XFree86 configuration keeps its information in a file called XF86Config, which is usually to be found in the /etc directory. When you're just getting started with UNIX, finding your way around the various directories can be confusing. A number of them have been set aside, by convention, for more or less specific uses, but the catch is that the set of directories and their uses tend to vary between different flavours of Unix. So, as for much else in this column, what follows isn't prescriptive, but is to help you get a handle on how it works.

- Expect to find the standard binary files intimately associated with the operating system itself in /bin.
- Devices (pseudo files which are actually connections to device drivers) are in /dev.
- /tmp is reserved for temporary files.
- The /usr directory will typically contain nothing but a number of other directories, like /usr/sbin, for binaries used for system administration, and /usr/bin, for frequently used binaries that aren't already in /bin.

The distinction between /bin and /usr/bin is fairly arbitrary. You might think that putting them under /usr implies they are the operating system enhancing binaries added by users, but these belong more properly (on many systems) in /usr/local/bin.

The /etc directory sounds as if it might be for everything else, but it, too, has a (roughly) defined role. These days the /etc location tends to be the place where all the configuration files are kept. But for historical reasons, /etc is also where you'll find networking utilities like ping and ifconfig. Oh, and the executable init scripts are probably in there, too. Yes, it's a bit of a mess, but it's a widely understood mess!

More recently, there's been a drive to formalise this directory tree more strictly, but in the meantime hold on to the idea that there is at least some rhyme and reason behind the layout. You get a feel for where the various files are lurking, and you can always supplement this with the system search utilities /bin/find and /usr/sbin/whereis. Or is that /usr/bin/find and /bin/whereis?

Sign of the times

I may not get the chance to reply to all your email (I do try) but feedback from you is definitely the making of this column, so please keep it coming.

There is, however, one kind of email I really don't like getting. No, I'm not talking about abusive email (rare, I'm glad to say) but the all-too-common mailing that expects me to do all the work and doesn't give me anything to go on.

Here's an example: "*I'm trying to connect my new Windows 95 PC to our IBM AIX database server for client-server access. The IBM is expecting a TCP/IP connection to service the client-server request. I've set up my PC with Microsoft's TCP/IP as one of its protocols (using Network Neighbourhood Add & Properties). What do I need to do now to establish a connection?*"

My first question is, who are you? The email isn't even signed. I had to fish about in the headers to discover the name, Ian MacDonald. What does it cost to spend a line or two introducing yourself? For example, "Hi, Chris, I'm a nuclear physicist at Cern, and we're using FreeBSD to control our particle accelerator. I happened to pick up a copy of *PCW*…"

Another point that is almost mandatory in mailings, not just to me but to newsgroups or whoever (and we talked a bit about this last month) is to indicate that you've gone some way towards trying to solve the problem yourself and you're not just leaning back, expecting someone else to do all the work. Even if you were, you'd need to define the problem rather better than Ian does here: "*The IBM is expecting a TCP/IP connection…*" doesn't help much. If there is such a connection and nothing's working, you have one set of problems. If you haven't managed to make the TCP/IP connection at all, that's a different problem.

The power of ping

The way you test whether a simple TCP/IP connection exists is to access the AIX machine's IP address from the client machine using the elementary utility called "ping" which should be available on any TCP/IP setup, and is certainly there in Windows 95. Now, I acknowledge there's a real problem for Windows 95 users: if they don't happen to know about the existence of ping (and why should they?) it's very hard to find out about it. Despite the extensive so-called "Help" Microsoft offers, there's no way (that I can see) for the complete novice to get from "network" to "ping". Contrast this with any decent UNIX operating system, where entering "apropos network" at the command line will return a list of associated commands: a little research among these will lead you to ping.

So ping may be all it takes to diagnose Ian's problem, or it may not. Who could possibly tell on the basis of the skeletal information provided?

SCO — sorted

By way of contrast, here's the sort of email I much prefer: "*I read your column regarding*

p268 ➢

Paul Rowlands' email mentions the HTML documentation that comes with SCO Open Server. Here's what it looks like. The help screen is a simple web browser (I've added Netscape 2.0.2 to my own Open Server setup, downloadable from the SCO web site) and the underlying HTTP server seems to be Apache, so this is more or less the same setup you get with Linux distributions these days

the problems you were having running SCO and DOS on the same (Apricot) system; seems like you had more trouble than I did … I have SCO OpenServer 5 (free), NT4, and DOS 6/Windows 3.11 all on the same system … My only problems are:
1. When I boot, I have to use fdisk to set either the Unix or the two Microsoft partitions active. The NT loader doesn't want to know about Unix.
2. SCO doesn't support my graphics card (a Number 9 Reality 332) so I have to run it in 640 x 480 mode — a bit grim, since it can support 1024 x 768 x 32K colours under NT and Windows…

"*Another gripe about SCO is that it seems impossible to feed back problems or get support. I know the product is free (but for a 'small handling charge'), but what is it there for? Is it to encourage me to buy it in my professional capacity?*

"*Have you noticed that SCO comes with an HTTP server already configured and running (dunno what particular flavour it is, but it contains O/S documentation)? Anyway, all the best, keep it up — we need a non-Microsoft perspective on the world.*"

Paul Rowlands

Paul also provided details of his hardware and made a few other useful points about SCO, which I've omitted for the sake of brevity. This is the kind of mail to which I can respond. I suppose the most serious point is the one about SCO's unapproachability. The good news is that this shouldn't be the case for readers of this column because Dave Gurr, marketing development manager at SCO UK, has been very helpful to me over setting up Open Server, and he's allowed me to publish his phone number and email address (*see "PCW Contacts", p269*) for any of you who need further help.

Please remember that the same rules (about which I've been writing) apply here, too: do your homework first. You'll find help about Open Server installation at www.sco.com, so please investigate here before contacting Dave.

As luck would have it, several other readers have written in with help for the problem about needing fdisk to switch between bootable partitions. My apologies for not having spotted this myself, but SCO has actually included some elementary dual boot software in the distribution. Robert Warner robert@softdesign.demon.co.uk discovered this when setting up a disaster recovery server that used two partitions, one for SCO and one for Windows 95.

He says: "*I installed SCO first, having the usual drivers problem because I was using a non-Adaptec SCSI controller…*"

(Minor technical carp, Robert. These aren't controllers, they're host bus adapters. With SCSI, the controller is embedded into the drive device itself.)

"*…Once I had that up and running, I proceeded to install Win95. This all worked fine except that when I booted the PC, it went straight into Win95 instead of getting the SCO boot prompt…*"

This is because Windows 95 assumes that it is the only operating system on the disk, and on install sets its own partition as the active partition. Robert fixed this using fdisk to make the SCO partition active. This gives you the SCO Unix boot prompt when you switch on, and if you just hit return at this point you'll be booted into Unix.

"*But,*" writes Robert, "*if you enter 'dos' and press return, it will boot from the Win95 partition. You get the best of both worlds.*"

That weasel word, "Open", again

A couple of months ago I quoted SCO's Michael Tilson, who was complaining that: "*The term 'open' has been hijacked. Because 'open' is good, everyone labels whatever they sell as 'open'.*"

It's a point I had raised with Caldera when it first came out with OpenLinux™ at the beginning of this year. But more recently, I found myself becoming increasingly irritated by this new development because of a supposed Caldera product called "OpenLinux Lite".

My initial objection to "OpenLinux" was that it capitalises on an established (non-proprietary) brand, but somehow manages to suggest that it is more "open" than Linux. Depending on your interpretation of the word "open" you might not accept this. A plain Linux implementation like the RedHat one, on the *PCW* CD a couple of months ago, can be shared and copied without restriction; Caldera OpenLinux can't. So which is the more "open"?

Well, the argument goes: Open with a capital "O" refers to adherence to standards and is not about whether or not you get the source code, or a reference to unrestricted redeployment. The new Caldera distribution is being aimed towards full UNIX compliance according to the specifications laid down by The Open Group. This is something most Linux users don't care about, so in this sense their plain old Linux is arguably less "Open" than Caldera's.

The OpenLinux Base product that currently sells for about £55 includes the Looking Glass desktop, the Netscape Navigator 2.02 web browser, Metro-X's proprietary accelerated X-Window system, the CrispLite text editor and other stuff. The version stripped of all this, and so freely distributable is, as I understand it, to be called "OpenLinux Lite". This seems to me to imply that it's somehow a cut-down version of Linux, which it isn't.

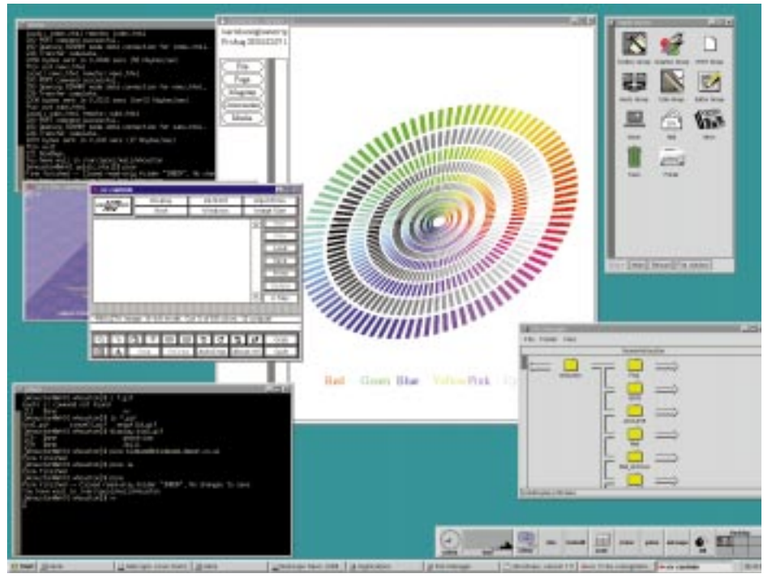

This screenshot, supplied by Martin Houston, organiser of the Linux branch of the UK UNIX User group (www.mh01.demon.co.uk) shows RedHat Linux running the latest fvwm97 window manager (the resemblance to windows from another well-known operating system is not accidental). You'll find a number of optional Linux interfaces at www.PLiG.org/xwinman/

Perhaps a more serious objection to OpenLinux Lite is that despite promises made at the beginning of this year, no mention of it has yet appeared on Caldera's web page and it doesn't seem to be obtainable (although you can download OpenDOS, Caldera's own version of DOS, bought in from Novell). I hope this will have changed by the time you read this. You can check www.caldera.com to find out, and if it still isn't there you can write to nancy.pomeroy@caldera.com to ask why not.

By way of contrast, I notice SCO has made good its promise to release a free single-user version of UnixWare on much the same basis as its currently very popular (to judge from your emails) and free Open Server. Find details of both operating systems at www.sco.com/offers/index.htm.

Why two operating systems? SCO is the largest commercial vendor of Intel-based UNIX systems, and when Novell's efforts to make a decent business out of UNIX fell apart a couple of years ago, SCO stepped in towards the end of '95 and bought UNIX from them. (At a bargain price, by the way: SCO needed to bring its own ageing version of UNIX up to date and paid Novell $60m for something that Novell had paid AT&T $320m for two years earlier!)

SCO's intention is to combine its Open Server with Novell's UnixWare, and the result of this will be appearing later this year as "Gemini". The free single-user version of UnixWare looks like something worth investigating. I've put in my own request for a copy, so expect news in this column soon.

More adventures with awk

As regular readers know, I like "awk". It's a good, simple, general-purpose language and much better than BASIC, in my view, as a way to get started in programming. It also comes free, in the shape of GNU's "gawk". The awk examples we've had in this column in the past have always required you to evoke gawk explicitly and add the program file and datafile as parameters, like this:

```
gawk -f progfile datafile
```

Recently, I decided I wanted to write an awk program as a utility that you can run directly as an executable in its own right. With shell files (written in csh, bash or whatever) the trick is to use what's called the "shebang notation". You do this by adding a magical first line to your script that tells the system where to find the language file that's going to run it.

```
#! /usr/bin/csh
```

The "#" (pronounced "hash" or "sharp") indicates that the line is a comment, and the immediately following "!" (or "bang") triggers the system magic. Most modern UNIX systems support this.

You can do the same thing with an awk script but what happens next isn't what you want: gawk loads and then tries to treat your script as if it were the data file. The trick (not immediately obvious from any of the documentation) is to use the -f indicator, just as you would at the command line.

Here's a skeletal example that also collects parameters from the command line:

```
#! /usr/local/bin/gawk -f

# How to pass command line params into gawk using ARGV, ARGC

# The ARGV array starts at 0, but that value will be "gawk"
# So we start at ARGV[1]

BEGIN {for (x = 1; x < ARGC; ++x)
            print ARGV[x]
            print "There were " ARGC - 1 " arguments,
not counting ARGV[0]"
     }
```

There's an amusing catch to this, though, if you're also trying to pass one or more data files into gawk from the command line, and I'll deal with that next month.

PCW Contacts

Chris Bidmead is a consultant and commentator on advanced technology. Email him at unix@pcw.co.uk.

SCO UK, Dave Gurr 01923 813548; email davidgu@sco.com