

YOU KNOW HOW Ken Thompson and Dennis Ritchie created Unix on a PDP-7 in 1969? Well, around 1971 they upgraded to a PDP-11 with a pair of hard drives.

The `/bin` vs. `/usr/bin` split (and all the others) is an artifact of this, a 1970s implementation detail that got carried forward for decades by bureaucrats who never question why they're doing things. It stopped making any sense before Linux was ever invented for multiple reasons:

packages, so let's add `/opt`! I'm still waiting for `/opt/local` to show up...

Of course, given 30 years to fester, this split made some interesting distro-specific rules show up and go away again, such as “`/tmp` is cleared between reboots, but `/usr/tmp` isn't.” On Ubuntu, `/usr/tmp`

Understanding the `bin`, `sbin`, `usr/bin`, `usr/sbin` Split

By ROB LANDLEY

When their root filesystem grew too big to fit on their tiny (half a megabyte) system disk, they let it leak into the larger but slower RK05 disk pack, which is where all the user and home directories lived and why the mount was called `/usr`. They replicated all the OS directories under the second disk (`/bin`, `/sbin`, `/lib`, `/tmp`...) and wrote files to those new directories because their original disk was out of space. When they got a second RK05 disk pack, they mounted it on `/home` and relocated all the user directories to this third disk so their OS could consume all the space on the first two disks and grow to *three whole megabytes*.

Of course they made rules about “when the system first boots, it has to come up enough to be able to mount the second disk on `/usr`, so don't put things like the mount command in `/usr/bin` or we'll have a chicken and egg problem bringing the system up.” The fact their tiny system disk was much faster than an RK05 disk pack worked in there too: moving files from `/bin` to `/usr/bin` had a significant performance impact on this particular PDP-11. Fairly straightforward, and also fairly specific to the hardware v6 Unix was developed on 40 years ago.

1. Early system bring-up is the province of `initrd` and `initramfs`, which deal with the “this file is needed before that file” issues. We already have a temporary system that boots the main system.
2. Shared libraries (introduced by the Berkeley guys) prevent you from independently upgrading the `/lib` and `/usr/bin` parts. Two partitions have to match or they won't work. This wasn't the case in 1974; back then they had a certain level of independence because everything was statically linked.
3. Cheap retail hard drives passed the 100 megabyte mark around 1990, and partition resizing software showed up somewhere around that time (partition magic 3.0 shipped in 1997).

Of course once the split existed, some people made other rules to justify it. Root was for the OS stuff you got from upstream and `/usr` was for your site-local files. Then `/` was for the stuff you got from AT&T and `/usr` was for the stuff that your distro, like IBM AIX or Dec Ultrix or SGI Irix, added to it, and `/usr/local` was for your specific installation files. Later, somebody decided `/usr/local` wasn't a good place to install new

packages, so let's add `/opt`! I'm still waiting for `/opt/local` to show up...
Of course, given 30 years to fester, this split made some interesting distro-specific rules show up and go away again, such as “`/tmp` is cleared between reboots, but `/usr/tmp` isn't.” On Ubuntu, `/usr/tmp`

doesn't exist, and on Gentoo, `/usr/tmp` is a symlink to `/var/tmp`, which now has the “not cleared between reboots” rule. Yes, all this predated `tmpfs`. It has to do with read-only root file systems. `/usr` is always going to be read-only in that case, and `/var` is where your writable space is. Moreover, `/` is mostly read-only except for bits of `/etc`, which they tried to move to `/var`, but symlinking `/etc` to `/var/etc` happens more often than not.
Standards bureaucracies, like the Linux Foundation (which consumed the Free Standards Group in its ever-growing accretion disk years ago), happily document and add to this sort of complexity without ever trying to understand why it was there in the first place. “Ken and Dennis leaked their OS into the equivalent of home because the root disk on the PDP-11 was too small” goes whoosh over their heads. ■

Rob Landley has been a geek since childhood, a Linux geek since 1998, and an embedded Linux geek since 2001.

Reprinted with permission of the original author. First appeared in *hn.my/bin* (busybox.net)

References:

- <http://cm.bell-labs.com/cm/cs/who/dmr/notes.html>
- <http://cm.bell-labs.com/cm/cs/who/dmr/hist.html>