

ZACK'S KERNEL NEWS

Maintainer Changes

Hans Verkuil posted a patch adding himself and Andy Walls as co-maintainers of the CX18 Video4Linux driver. Arnd Bergmann added a MAINTAINERS entry listing Jeremy Kerr as the SPU Filesystem official maintainer. Zhang Wei posted a patch adding Li Yang as his co-maintainer of the Freescale DMA driver. Steve Wise has listed Divy Le Ray as the CXGB3 Ethernet driver maintainer and listed himself as the CXGB3 iWarp RNIC driver. Paul Jackson sent in a patch announcing that Simon Derr would no longer be his co-maintainer of the CPUSets code and was being replaced by Paul Menage. Jean Delvare announced that Ben Dooks would be joining him as co-maintainer of the I2C subsystem. Auke Kok stepped down from being the primary maintainer of the Intel Ethernet drivers and handed maintainership to one of his co-maintainers, Jeff Kirsher. Jesse Brandeburg will remain a secondary maintainer on that project. Finally, Alan Cox has made a few more alphabetization fixes to the MAINTAINERSHIP.

The Linux kernel mailing list comprises the core of Linux development activities. Traffic volumes are immense, often reaching ten thousand messages in a given week, and keeping up to date with the entire scope of development is a virtually impossible task for one person. One of the few brave souls to take on this task is Zack Brown. Our regular monthly column keeps you abreast of the latest discussions and decisions, selected and summarized by Zack. Zack has been publishing a weekly online digest, the Kernel Traffic newsletter for over five years now. Even reading Kernel Traffic alone can be a time consuming task. Linux Magazine now provides you with the quintessence of Linux Kernel activities, straight from the horse's mouth.



Porting Linux to ... Everywhere?

Octavian Purdila, Stefania Costache, and Lucian Adrian Grijincu are essentially trying to port Linux to run on all other programming projects. They call it the Linux Kernel Library project, and the goal is to convert various things, such as the Linux Virtual File System, into generic libraries that would allow anyone to insert them into their projects.

If successful, any operating system would natively be able to support any Linux feature, just by using this library. Their approach is to create the library as a direct port of Linux to a virtual architecture that could then be used by anything that wanted it.

Octavian, Stefania, and Lucian have made some progress already and are looking for volunteers to help keep the port up to date with the kernel.

GPL Violations

Every once in a while, someone tries to get around the GPL. Jo Shields recently noticed that the Taiwanese company Compro Technology was offering a binary-only driver for their product, only it turned out that the driver was actually a complete kernel image, a modified version 2.6.17.

Jo contacted the company as well as various members of the press, and Compro took the code down.

On their website, the company explained that they were in the process of complying with the GPL and would have their drivers ready in a few months. Bernd Petrovitsch wondered aloud whether they were perhaps just waiting for folks to forget about the violation before rereleasing it in another binary-only form. Time will tell.

2.4.36.4 Released; Users Encouraged To Upgrade

Willy Tarreau has released 2.4.36.4, which includes some important security fixes. Alexander Viro had identified and fixed an SMP problem as well as a *dnotify* race condition.

UBIFS Official Tree Inclusion

Artem Bityutskiy made a formal request to have UBIFS, a relatively new Flash filesystem, merged into the mainline tree during the next merge window.

Artem asked Stephen Rothwell to include his git tree in the Linux-Next tree, aiming for submission to Linus Torvalds at the next opportunity.

Stephen had no objection to this. In general, Stephen's requirements for inclusion in Linux-Next are that code be reviewed and tested as well as can be managed and that the submitter intends to have the code accepted in the next merge window. Because Artem had done all that, Stephen was happy to include Artem's tree.

Andrew Morton also approved, although he hadn't looked at any of the code himself.

Andrew knew that Christoph Hellwig had made a number of technical suggestions and wanted to know the status of Artem's work in response to that. Artem said that he and the other UBIFS developers were actively addressing Christoph's concerns.

Regarding Christoph's suggestions, Artem did say that NFS support would probably not happen in the near future, but that some ideas were already percolating as to how to fix that in the long term. None of this seemed to pose a significant obstacle, and it looks as though UBIFS will be merged (or at least will go to Linus for consideration) for the 2.6.27 release.

Purging OSS

The final PCI OSS driver has been removed from the Linux tree. In the old days, Muli Ben-Yehuda had resisted taking out the Trident 4DWave/SIS 7018 PCI Audio Core, but he's moved on to other projects and no longer has the hardware to maintain the driver. With his permission, and given that an ALSA driver that supports the same hardware exists, Adrian Bunk submitted the patch to remove it.

Status of the Big Kernel Lock

Linus Torvalds posted a patch to undo a large change to the kernel's locking system. The big kernel lock (BKL) had been changed from a spinlock, with some latency issues, to a semaphore. This had been a big win for folks doing work supporting patient monitors in hospitals (or playing first-person shooters).

The semaphore solution caused huge performance problems under certain benchmarks, and the solutions to these performance issues still left lingering problems. So Linus decided that the simple way out was to go back to the way things had been before the change.

Linus also indicated that the only way around the spinlock approach would be to eliminate the BKL entirely. Ingo Molnar, a self-proclaimed "latency junkie," was not pleased with this, but he accepted Linus's premise that the solution was to pull out the BKL at the root. It turns out ditching the BKL is hard! It took people with the raw talent of Alan

Cox (quoth Ingo) to wind their way through the maze of semantics in order to know what to change and how to change it. According to Ingo's calculations, even with folks like Alan already hard at work on the task, at the current pace, it would be more than 10 years before the BKL was gone. Ingo explained the problem, saying, "... its dependencies are largely unknown and invisible, and it is all lost in the haze of the past ~ 15 years of code changes. All this has built up to a kind of Fear, Uncertainty, and Doubt about the BKL: nobody really knows it, nobody really dares to touch it, and code can break silently and subtly if BKL locking is wrong."

Ingo's solution was to create a git tree specifically to rein in some of the more horrifying BKL oddities, with the goal of ultimately making the BKL easier to remove en masse. One of the main first steps was to extract it from the core kernel code and move all its ugliness into

something that could one day theoretically be replaced on its own, changing the behavior throughout the kernel once an improved implementation could be designed. As Ingo put it, "Once this tree stabilizes, elimination of the BKL can be done the usual and well-known way of eliminating big locks - by pushing it down into subsystems and replacing it with subsystem locks, and splitting those locks and eliminating them. We've done this countless times in the past and there are lots of capable developers who can attack such problems."

Andi Kleen liked the plan, but figured why wait? Instead of having a separate git tree for all the changes, just do them in the official branch. Linus was also pleased to see Ingo working on this and had some suggestions about how to organize the git tree. A bunch of other high-powered hackers also turned their attention to the various technical issues involved.

Linux Magazine Exclusive

