

ZACK'S KERNEL NEWS

Maintainership Updates

Alessandro Rubini has abandoned maintainership of the “Mouse and misc devices (general)” code, which actually hasn’t existed for a long time. The generic mouse interface has been part of the input subsystem maintained by Dmitry Torokhov, and miscellaneous drivers are handled individually. Both the maintainer and the project itself were out of date in the MAINTAINERS file, and Jiri Kosina posted a patch to remove it.

Rodolfo Giometti has added an entry to the MAINTAINERS file, listing himself as the official PPS (pulse per second) maintainer. PPS provides an API for precise time measurements that can be used to synchronize clocks on disparate computers. Although several implementations are available for PPS in Linux, this one is Rodolfo’s. Its benefits are that it makes it easy to add and manage new sources of timing information into a pool, with a consistent interface for each source.

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.



Kristoffer Ericson created a MAINTAINERS entry for the Epson S1D13XXXFB driver and listed himself as the official maintainer, “... since I have 2 platforms currently using it,” he said. This controller is used for various LCD screens, typically embedded as display panels in other devices such as printers, car dashboards, and whatnot.

Theodore (Ted) Y. Ts’o has updated the MAINTAINERS file, listing himself as the ext4 filesystem maintainer instead of Andrew Morton and Stephen Tweedie, although this doesn’t represent an actual shift in maintainership. According to Ted, the entry was originally just copied from the ext3 entry and never had been accurate. Ted’s patch just corrects the inaccuracy.

An interesting detail is that, although Andreas Dilger is not included in his patch as a co-maintainer, the list of appropriate email addresses for patch submissions, bug reports, and so forth includes Andreas’ email address alongside Ted’s. Andreas’ email was also listed alongside Andrew’s and Stephen’s in the original mistaken entry. Typically, this would mean that he puts a lot of work into the project and can act as a lieutenant to Ted, in the same way that Linus Torvalds has trusted lieutenants who have an easier time getting code into his tree than other folks.

Len Brown posted a patch removing Kristen Carlson Accardi as the maintainer of the Docking Station driver and listing Shaohua Li as the replacement. Both Kristen and Shaohua work for Intel, so we can assume some internal shuffling is going on there. Kristen had been the official maintainer since July 2006, according to another patch Len posted at that time.

Thomas Gleixner has given maintainership of rt-tests to Clark Williams. The rt-tests code contains various test programs for the kernel’s real-time support. Cyclctest, the original test program that spawned rt-tests, is still included in the package and tests the interrupt/wakeup/

scheduler/userspace chain of events within the running kernel. In fact, the whole package has been so useful that it is now included in the official kernel tree, which is something Thomas had never envisioned when he first created Cyclctest.

Kentaro Takeda has posted a patch listing himself and Tetsuo Handa as co-maintainers of the Tomoyo security module. This GPL’d project was developed by the NTT Data Corporation and implements MAC (mandatory access control) security policies on files, networks, and other system resources.

Peter Oruba removed himself as maintainer of AMD Microcode Update Support and listed Andreas Herrmann instead. This code handles the insertion of new updates into the microcode that controls how the AMD CPU functions. In July 2008, Peter originally introduced this code into the kernel and has maintained it since then.

Mike Frysinger posted a patch to the MAINTAINERS file listing himself and Subrata Modak as co-maintainers of LTP (Linux Test Project). This project, sponsored by SGI and IBM, provides a suite of about 3,000 tests addressing reliability, robustness, and stability of the Linux operating system.

Phillip Lougher has added an entry to the MAINTAINERS file, listing himself as the official SquashFS maintainer. SquashFS is a read-only compressed filesystem, intended for embedded systems and other small Linux utilizations – or really anywhere else. Currently it supports gzip compression because that’s the only form of compression available natively in the Linux kernel, but patches are available at <http://www.squashfs-lzma.org/> for anyone who wants to try LZMA compression.

Michal Januszewski added an entry to the MAINTAINERS file for the uvesafb driver, listing himself as the official maintainer. Uvesafb is a generic framebuffer driver with a variable refresh rate that replaces vesafb-tng.

Big Lock Eradication

Frederic Weisbecker has created the Big Kernel Lock Tracer, which tracks the time penalty incurred by the big kernel lock (BKL). This could help identify specific areas in which the BKL should be removed before others. Linus Torvalds has pushed to get rid of the BKL and replace it with simpler locking structures. However, not all situations can be resolved with a generic lock, and the locking code has been scattered throughout the kernel, making any kind of concerted effort problematic. One attempt has been to “push” all BKL code into its own part of the source tree to centralized it so it can potentially be dealt with all at once. Developers across the spectrum of kernel development have been engaging in a search and destroy mission against the BKL for months now, but a lot of work remains. Frederic’s tool might help developers focus their attention on the areas that would benefit the most.

Kernel Version Numbering

Kernel version numbers might be redefined soon, which recently inspired debate between proponents of changing the system and those who wanted to leave the system alone.





Greg Kroah-Hartman, the main advocate of changing the current system, suggested using year numbers, followed by the number of each release done in that year, followed by the number of stable versions of that release. Thus, 2009.3.5 would be the fifth stable release of the third version released in 2009. This, he said, would prevent the version numbers from getting too large (e.g., 2.6.25).

Given the opposition to the change among big-time developers and the plethora of well-meaning suggestions, the system probably won’t change anytime soon. However, the original complaint came from Linus Torvalds, and Greg started the discussion as a way to come up with a recommendation for

him. With Linus behind the change, it seems as though some kind of new versioning scheme will eventually get through, although the discussion might make Linus think twice.

Redefining Load Average

Sena Seneviratne and David Levy have been cooking up some changes to the load average calculation in Linux. Their idea is that instead of just calculating a single number representing the entire load average on the system, multiple statistics should be available. For one, they want to separate disk I/O load from the rest of the load average calculation so it can be viewed independently. Also, they want to enable per-user load average calculation, so each user could see their effect on the system. Arjan van de Ven exhorted them to add statistics only and not to change the behavior of any numbers that the system already produces.



Linux Magazine Exclusive