

ZACK'S KERNEL NEWS

Mea Culpa

Before anything else, I'd like to correct the record. Last month, I referred to Rodolfo Giometti as Rudolfo. Sorry about that! And thanks for the correction, Rodolfo!

An Attempt at a New Userspace Block Device

Zachary Amsden has introduced support for userspace block devices, similar to FUSE (Filesystem in Userspace) and CUSE (Character Device in Userspace). Zachary calls his aBUSE. As Alan Cox pointed out, Zachary's code does something similar to the NBD (Network Block Device) project, which has existed for a while. But Zachary said he already knew about NBD and felt that aBUSE was sufficiently different to justify itself. For one thing, NBD requires access to a socket, which could lead to deadlocks under some circumstances, and NBD is just more complex in general.

Tejun Heo agreed that the similarity to NBD was a big drawback and also pointed out that FUSE could be used to accomplish much the same thing as aBUSE as well; however, as Zachary remarked, such an implementation would be very annoying to attempt.

Nevertheless, the comparison with NBD ultimately convinced Zachary, and he decided to rely on NBD for his own needs and leave aBUSE in the mailing list archives in case anyone felt the need to explore it further.

ARM Mailing Lists Forked over Open Posting

After a big flamewar, Russell King quit maintaining the ARM mailing lists. The dispute was primarily over whether to allow non-subscribers to post to the list. Traditionally, in spite of problems with spam, Linux mailing lists remain open so as to make it as easy as possible for regular users to produce and submit bug reports. But not everyone agrees. Russell, maintaining the ARM lists himself, had found the spam problem too difficult to deal with, without restricting posting privileges to subscribers only. When Pavel Machek complained about this one time too many, Russell closed the lists down completely.

David S. Miller, Alan Cox, and Theodore Y. Ts'o all disapproved of Russell's hard-line approach to the debate. At the same time, folks like Bartłomiej Zolnierkiewicz complained that people should give Russell a break, especially people who were not involved in ARM development. But, meanwhile, Pavel submitted a patch migrating 19 ARM mailing lists to vger, which had more open posting policies. And, Joe Perches sent out a similar patch, covering more than 50 mailing lists.

The sudden transition to vger caused a bit of confusion among ARM developers. Martin Bogomolni remarked, "This confusion, between shutting down the list, and over which list should take precedence could not have come at a worse time. I have a number of Cortex-A9 re-

lated patches in the pipeline to add a new board, platform, and power management features."

Meanwhile, with Russell's approval, David Woodhouse set up a new batch of ARM mailing lists at lists.infradead.org and migrated all the subscribers from the old lists to the new ones. This caused more confusion when it turned out that a lot of people had set their subscriptions to receive no email, so they'd be able to post to the mailing list in spite of the subscriber-only restriction. David W. had not carried over that configuration detail, and when folks started complaining about their mail clients being flooded with ARM mail all of a sudden, his response was that, yes, he should have carried over the configuration, but that, "Still, it's done now – people can easily unsubscribe. Apologies to those inconvenienced, but it's probably better that way than for people to drop off the list without noticing."

Meanwhile, developers were also confused about whether they should subscribe to the lists on vger or the lists on infradead; the confusion was worsened by the fact that some of Russell's lists had been merged together into single lists during the transition.

The whole thing will undoubtedly be resolved soon; but, if the goal was to upset the most amount of people during a mailing list transition, that was probably achieved.

The Linux kernel mailing list comprises the core of Linux development activities. Traffic volumes are immense, often reaching 10,000 messages in a 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.



kernel.org Homepage Revamp

Periodically, people complain about the scarcity of links to various key kernel trees on kernel.org. Recently, a bunch of folks asked why the -mm and linux-next trees were not included on the kernel.org homepage. John Hawley, one of the kernel.org maintainers, said, "the way kernel trees are placed in the hierarchy has changed from ... when the code was first put together. ... basically [it's] a complete re-write of the code base at this point." A few weeks later, he announced, "I've pushed out an update that should incorporate the expected trees now; this does eliminate the 2.2 and all but the last 2.4 tree (2.4.37.5), but does include all of the stable 2.6.x trees, the snapshots, and linux-next."

Andrew Morton's -mm tree still didn't show up on the main page, and he remarked, "I need to find a way to shove most of -mm into linux-next." So the -mm tree might ultimately be going away in favor of linux-next, and it might not get onto the kernel.org homepage with the rest of the primary trees.

ReiserFS Locking; How Much Testing Do Patches Need?

Frederic Weisbecker has been working on getting the Big Kernel Lock (BKL) out of ReiserFS. Originally, the ReiserFS changes were part of the overall lot of anti-BKL patches going into the official git tree, but they soon migrated to their own ReiserFS-3-specific tree. The locking mechanism that Frederic has settled on is based on a recursive mutex and is mostly faster than the BKL. In at least one benchmark, however, it performed less well.

An interesting discussion emerged out of Frederic's patches. Andi Kleen pointed out that with all the benchmarking, there were no stress tests included to make sure the code actually behaved appropriately and didn't corrupt data on disk. Especially in the case of filesystems, data integrity tends to be considered sacrosanct. But Ingo Molnar argued that beyond basic due diligence, it wasn't actually necessary to knock oneself out testing patches before submission. Running the gauntlet through linux-next and the -mm tree would produce a far wider range of test-

ing environments than anything a contributor might have available.

Of course, nothing's wrong with using stress test software when it's available. Chris Mason pointed out that the testing script *stress.sh* would put ReiserFS through its paces, and Frederic was very happy to do that before sending in the code for further testing.

This seems to be an extension of the new "include early, include often" philosophy indoctrinated when Linus Torvalds decided to discontinue the whole even/odd development style, wherein even-numbered trees like 2.6 would focus on stability, and odd-numbered trees, like the yet-to-be-seen 2.7 tree, would focus on development. Since that decision, the effort to ensure kernel stability has been pushed into a number of different areas, including the 2.6.x.y stable tree that comes out for each official release; the linux-next, -mm, and other specialized trees; and the Linux operating system distributions themselves.

Strategy for Merging IIO

Developer Jonathan Cameron wanted to get the IIO (Industrial Input/Output) subsystem into the main Linux kernel tree. Cameron asked if getting the IIO code into the staging tree was the best way to go about moving it to the main tree.

The IIO subsystem provides a collection of features for supporting hardware sensors such as gyroscopes, light sensors, and other tools that have a relatively high update frequency.

Greg Kroah-Hartman had no problem accepting the IIO code into the staging tree, but he wanted to see a to-do list that lays out a clear path toward migrating the IIO subsystem from staging to the main kernel code areas.

If all goes well with the promised to-do list and the subsequent discussion, it looks like the Industrial Input/Output subsystem is heading into the main tree sometime in the near future. Anyone porting Linux to the Segway will be happy to hear this.

THE MATHEMATICS OF HUMOUR

TWELVE Quirky Humans,
TWO Lovecraftian Horrors,
ONE Acerbic A.I.,
ONE Fluffy Ball of Innocence and
TEN Years of Archives
EQUALS
ONE Daily Cartoon that Covers the
 Geek Gestalt from zero to infinity!

Over Two Million Geeks around the world can't be wrong!
 COME JOIN THE INSANITY!



UserFriendly.org