Zack's Kernel News

Chronicler Zack Brown reports on the latest news, views, dilemmas, and developments within the Linux kernel community.

By Zack Brown

Andre Hedrick: In Memoriam

This edition of Zack's Kernel News is dedicated to Andre Hedrick, a prolific kernel hacker who committed suicide in July. He did tons of kernel work for many years on the IDE driver, and later the ATA subsystem it morphed into.

I remember him as being one of the more argumentative of kernel folks. He always seemed to be explaining why something was absolutely required or something was absolutely impossible because of the convoluted insanity of IDE hardware. The problem often was: There are hardware standards, and then there is reality. Often the hardware manufacturers didn't quite conform to the standards, but Linux had to support them anyway. Sometimes the standards themselves had missed some crucial nuance. Andre undoubtedly found it frustrating to debate such issues with people who he felt didn't fully understand the hardware under discussion.

Sometimes his enthusiasm to help users put him on the losing side of an argument. At one point in 2000, he found that if someone hacked root, they could brick the disks on the system, making them completely unusable, but no one wanted to take his patch! They all said it was pointless to try to protect the system from a malicious root user and that there were plenty of other nasty things root could do if it wanted. This is actually standard policy for Linux development: Presumably, if someone has root, they also have the ability to hit the hardware with a hammer. Root already has the keys to the kingdom, so why bother nit picking?

But Andre didn't see it that way. In this particular case, he probably felt that if someone hacked into root from outside, they at least shouldn't be able to do actual hardware damage. He also found it extremely frustrating that no one seemed to take him seriously. I remember him saying repeatedly in different email messages during the conversation, "Here is your SECURITY HOLE! JOE-SIX-PACK-HACKER can fry your butt."

Notwithstanding the various flame wars, Andre wanted to help people, and he dedicated quite a significant portion of his life to one of the thorniest, ugliest, most fathomless parts of the kernel – getting all the disks in the world to work right. Thanks for doing that for us, Andre.

Simplifying Config For Regular Users

Once in a while Linus Torvalds throws a feature idea out to the mailing list in the hope that someone will pick it up and implement it. Some ideas are more difficult to implement than others. Ultimately, no one was able to give him the revision control system he wanted, and he ended up having to write one himself.

This time, he was concerned about kernel configuration options. Specifically, he said the various Linux distributions (Red Hat, Debian, etc.) tended to have complicated kernel configuration requirements that went well beyond simply identifying the particular hardware installed on the system and that those requirements would change over time, making the problem even thornier. He said, "For example, F14 (iirc) started using TMPFS and TMPFS_POSIX_ACL/XATTR for /dev. And starting in F16, the initrd setup requires DEVTMPFS and DEVTMPFS_MOUNT. There's been several times when I started with my old minimal config, and the resulting kernel would boot, but something wouldn't quite work right, and it can be very subtle indeed."

Linus's idea was to have a set of distribution-specific and version-specific configuration options so users could initiate the configuration process by identifying the minimal setup required for their system. Once they'd identified the particular version of their distribution, they could continue the configuration process, selecting and disabling features as they saw fit.

The whole point, Linus said, was to make it possible for ordinary users to compile their own kernel without fear of messing things up and being left with subtle – or extreme – breakage. He also pointed out that simply copying the config file that came with the distribution – or doing a `make localmodconfig` – was not a good solution because it included a number of options that weren't really required.

A lot of discussion resulted in response to this. Dave Jones of Red Hat pointed out that this would have to be an ongoing process because, in some cases, a dependency catches even the developers themselves by surprise.

ZACK BROWN

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**.

He said, "We recently found out our virtualisation guys started using sch_htb for example, and we inadvertently broke it when we moved its module to a 'not always installed' kernel subpackage."

Linus was on board with the ongoing process idea. He replied to Dave, saying that even an educated guess would be an improvement. He pointed out that he really wanted two main things. First, each option selected by the distribution config should have a comment explaining why it's needed, and second, there should be a real effort made to exclude configuration options that aren't really needed. Linus said, "Other than that, even if it only gets you *closer* to a kernel that works with that distro, I think it doesn't have to be all that perfect. Because the alternative is what we have now."

Casey Schaufler also responded to Linus's initial post, objecting that the distribution configurations might tend to favor particular projects like SELinux, making that feature required, and excluding competing projects like LSM from consideration.

Linus agreed that this would definitely happen, but he said it didn't matter. The whole point of distribution-specific configurations would be to allow ordinary users to build a kernel easily. Anyone experimenting with special features and complex situations, was already going to be a much more advanced kind of user and wouldn't be using the distribution-specific configurations in the first place, Linus said.

The "SELinux versus LSM" debate didn't die right away, though. David Lang protested that users who set up a distribution configuration would be locked into certain features because they would be defined as dependencies rather than just desired features. There'd be no way to disable them. But Linus (and others) pointed out that even in David's worst-case scenario, it was still possible to edit the config file by hand, remove the distribution-specific selection, keep all the features that the distribution depended on, and then go through the `menuconfig` process again, unencumbered by the restrictions imposed by the distribution.

Linus's point was that the group of users who would be availing themselves of this new feature would not care whether SELinux or some other security system was the one being used, and the people who *did* care, would be more than capable of getting the config file they wanted.

A number of folks continued debating the ins and outs of how to implement the feature and what sort of caveats to watch out for. Clearly, however the final version ends up looking, it will soon be much easier to compile a kernel for a particular distribution, without introducing subtle breakage. Some sort of accommodation probably will be made to people who want to experiment with new or upstart features, without editing config files by hand.

Written-for-Linux-Magazine-by:

Alexandre Pereira da Silva wanted to add a "Tested-by:" signature to Git commits so testers could be CC'd on follow-ups. Joe Perches had no objection to that, but Andrew Morton said that "testers" were often just regular users somewhere in the world who'd reported a problem and who had no interest in doing further tests or in reading technical discussions about the problem.

Andrew also pointed out that changelog signatures varied quite a bit. He did a quick data-mining exercise and found a number of signatures that people had just made up and that no one had caught before the patches had made it into the kernel. Some were About-fscking-timed-by:, Antagonized-by:, and Fatfingered-by:, not to mention Sort-Of-Acked-By:. Tons of others were found, many of them pretty funny.

Andrew was not a huge fan of these tags, though. He said, "This fashion of adding new and innovative changelog tags just creates inaccuracy and work for people who want to mine that data for something useful, as Alexandre is finding out."

Joe told Andrew to relax and pointed out that the actual percentage of creative tags was quite low. He did his own data mining, which showed that virtually all the crazy tags were just one-offs. He added that he himself was responsible for AOLed-by:, Blame-taken-by:, Heckled-for-on-IRC-by:, and iSigned-off-by:.

Probably the fun will continue, to the consternation of data re-users everywhere. ■■■