

# ZACK'S KERNEL NEWS

## New Kernel Tester Guide

Michal Piotrowski announced some new documentation, in the form of the "Linux Kernel Tester's Guide," freshly translated into English by Rafael J. Wysocki [1].

A bunch of folks cheered at this development, and Jan Engelhardt asked whether there was an HTML version available as well.

At almost 70 pages, this document is written like a book and covers a huge amount of material. The authors have put a lot of time and effort into this, and it shows.

## Detecting ext4 Corruption

Mingming Cao added code to ext4 to store checksums of the filesystem journal to detect any corruption as soon as possible. After some technical comments and a few criticisms by folks like Andrew Morton, Girish Shilamkar posted an updated patch.

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.



## Distributed Replicated Disk Support Rejected

Lars Ellenberg and Philipp Reisner submitted their DRBD (distributed replicated block device) project code for inclusion in the official kernel. Lars made the case fairly well, saying the code was stable, had an existing user base, and would be discussed at upcoming conferences such as linux-conf.eu.

The response was not universal acclaim. Jesper Juhl noticed a variety of style and whitespace problems, and Lars confirmed that *scripts/checkpatch.pl* reported more than 2000 style complaints. He proceeded to wade through and fix them. Andi Kleen and Jens Axboe also pointed out that Lars had not posted all the code to the list but had just linked to a git repository. To get folks looking at the code, they recommended posting all the patches directly to the mailing list for easy viewing.

Sam Ravnborg, looking at the code himself, noticed some strange *typedef* and *volatile* usages in addition to odd macros. But he also suggested finishing off the style and whitespace fixes before getting into a huge discussion about those details. In a later email, Sam also suggested giving a better description of what the feature actually did.

Lars ran along as fast as he could, fixing problems as folks pointed them out. It turned out that some folks' comments would be very difficult to fix, particularly comments involving redesigning portions of the code. Over the course of the discussion on the linux-kernel mailing list, it became clear that DRBD would not be ready for acceptance into the main kernel tree for some time. But all the criticisms, though in some cases large, were friendly; and all the responses, though in some cases discouraging, were quick and relevant.

It seems clear that Lars and Philipp intend to address all concerns and come back with a new version based on everything they've learned from the discussion. To that extent, DRBD seems likely to make it in eventually.

## SDIO Flash Port Support

Pierre Ossman announced that SDIO (secure digital input/output) support would be merged into the mainline kernel soon. SDIO provides a way to use SD Flash Memory ports for devices other than Flash Memory cards. Available devices fitting the SD form factor include Bluetooth adaptors, modems, barcode readers, and FM radio tuners. Pierre and Nicolas Pitre worked together to code up the entire feature from scratch. Pierre also said that while basic functionality had already been implemented, there were still many performance optimizations yet to be coded.

## Hilscher CIF DeviceNet and Profibus Cards Support

Greg Kroah-Hartman and Hans-Jürgen Koch signed off on a patch by Hans J. Koch to support Hilscher CIF DeviceNet and Profibus cards. The kernel driver was small and just provided a gateway for a user-space application that handled all the heavy-duty work. Jan Engelhardt thought the code was ugly, Greg agreed, and Hans-Jürgen explained that the author hadn't paid much attention to kernel coding style or requirements for getting code into the kernel. Hans-Jürgen pointed out that the larger portion could be extracted into user space with no great problem. This, he said, was largely thanks to UIO going into the kernel recently, providing a generic way for drivers to put portions of their features in user space. Most of the driver's ugliness is outside of the kernel. Still, the code may be just ugly enough to keep it out of the main tree anyway.

## Kernel Archaeology

Interest among long-time kernel developers is becoming serious about getting a complete git repository of the entire Linux kernel history. Tracking down all the patches and release announcements is a huge task. If someone does put together such a repository, Linus Torvalds say he will comment as many of the patches from memory as he can.

## Maintainership

Anton Vorontsov posted a patch listing himself as the “POWER SUPPLY CLASS/SUBSYSTEM and DRIVERS” maintainer along with David Woodhouse.

Mike Sharkey, of Pike Aerospace Research Corporation, offered to maintain the currently unmaintained parallel port driver because his company is a heavy user of those devices. He asked what outstanding issues exist in that driver that he might start to address. Randy Dunlap gave him a link to <http://bugzilla.kernel.org/> that showed some bug reports related to the parport driver.

Valerie Henson posted a patch to remove herself as the official maintainer of the Tulip driver, and as she put it, “let her roam free, FREE!” Jeff Garzik applied the patch with no further comment. This happens just a few months after Valerie quit her job at Intel to become a Linux consultant. Wim Van Sebroeck posted a patch that, among other things, listed Mike Fryssinger as the official maintainer of the Blackfin Watchdog driver.

## The Container Convention

Serge E. Hallyn announced a small meta-convention that will be held during LCE (LinuxConf Europe) 2007 and focus on containers in the kernel. The meeting’s goal would be to grab a bunch of folks during the main convention, probably on September 3 or 4, and talk about current containers work over the course of a few half-hour presentations. To ensure maximum attendance, Serge added that they are preparing to set up phone lines into the conference meeting room, for anyone who can’t make it to the actual conference this year.

## Supporting Legacy Systems

Jonathan Campbell wrote some code to shrink the kernel on legacy hardware by not compiling various pieces of unneeded code into the binary. Jonathan noticed that even when compiling for the 386 processor, the kernel still included support for CPUID, the Pentium TSC register, DMI information, and other stuff not needed on such an old system.

After removing a bunch of these things, Jonathan found he could create a 450KB binary. There was a bit of discussion about his patch, but no one was blown away by the prospect of being able to better support old hardware. Some interesting points did come out of the conversation, though. For instance, creating a smaller binary has value because if the whole kernel can’t be loaded into RAM, the system can’t boot.

Once the system is up and running, however, the size of the binary doesn’t matter nearly as much as how much memory the running kernel uses, because all of user space must make do with whatever memory remains.

This is the sort of patch that goes in largely unnoticed, but to the ecstatic joy of a small number of people.

### INFO

[1] Linux Kernel Tester’s Guide: <http://www.stardust.webpages.pl/files/handbook/handbook-en-0.3-rc1.pdf>

Advertisement